

2nd
EDITION

DAVID ARCHER

GLOBAL WARMING

Understanding the Forecast



VISIT...

LANZAROTE
Caliente.COM

Global Warming

Understanding the Forecast

Second Edition

David Archer

d-archer@uchicago.edu
University of Chicago



WILEY

JOHN WILEY & SONS, INC.

*Dedicated to George Lockwood
and the spirit of curiosity*

Vice-President & Publisher:	Jay O'Callaghan
Executive Editor:	Ryan Flahive
Editorial Assistant:	Darnell Sessoms
Marketing Manager:	Margaret Barrett
Senior Media Editor:	Lynn Pearlman
Media Specialist:	Aaron Gass
Photo Researcher:	Jennifer MacMillan
Senior Production Manager:	Janis Soo
Production Manager:	Joyce Poh
Assistant Production Editor:	Annabelle Ang-Bok
Designer:	Alson Low
Cover background photo:	Paul Souders/Aurora Photos, Inc.
Cover top photo:	James Balog/Aurora Photos, Inc.
Cover center photo:	Science Source/Photo Researchers
Cover bottom photo:	Peter Essick/Aurora Photos, Inc.

This book was set in 10/12 Times Roman by Laserwords Private Limited and printed and bound by RR Donnelly (Von Hoffman).

This book is printed on acid free paper.

Founded in 1807, John Wiley & Sons, Inc. has been a valued source of knowledge and understanding for more than 200 years, helping people around the world meet their needs and fulfill their aspirations. Our company is built on a foundation of principles that include responsibility to the communities we serve and where we live and work. In 2008, we launched a Corporate Citizenship Initiative, a global effort to address the environmental, social, economic, and ethical challenges we face in our business. Among the issues we are addressing are carbon impact, paper specifications and procurement, ethical conduct within our business and among our vendors, and community and charitable support. For more information, please visit our website: www.wiley.com/go/citizenship.

Copyright © 2012, 2007 John Wiley & Sons, Inc. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc. 222 Rosewood Drive, Danvers, MA 01923, website www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030-5774, (201)748-6011, fax (201)748-6008, website <http://www.wiley.com/go/permissions>.

Evaluation copies are provided to qualified academics and professionals for review purposes only, for use in their courses during the next academic year. These copies are licensed and may not be sold or transferred to a third party. Upon completion of the review period, please return the evaluation copy to Wiley. Return instructions and a free of charge return mailing label are available at www.wiley.com/go/returnlabel. If you have chosen to adopt this textbook for use in your course, please accept this book as your complimentary desk copy. Outside of the United States, please contact your local sales representative.

Library of Congress Cataloging-in-Publication Data

Archer, David, 1960–

Global warming : understanding the forecast / David Archer. – 2nd ed.

p. cm.

Includes index.

ISBN 978-0-470-94341-0 (pbk.)

1. Global warming. 2. Global temperature changes. 3. Greenhouse effect, Atmospheric. 4. Global warming–Political aspects. 5. Global warming–Economic aspects. I. Title.

QC981.8.G56A73 2012

551.6–dc22

2011015260

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Contents

<i>Online Models</i>	v
<i>Preface</i>	vii
1 Humankind and Climate	1
Part I The Greenhouse Effect	7
2 Blackbody Radiation	9
3 The Layer Model	19
4 Greenhouse Gases	29
5 What Holds the Atmosphere Up?	43
6 Weather and Climate	57
7 Feedbacks	73
Part II The Carbon Cycle	87
8 Carbon on Earth	89
9 Fossil Fuels and Energy	103
10 The Perturbed Carbon Cycle	119
Part III The Forecast	133
11 The Smoking Gun	135
12 Potential Climate Impacts	153
13 Decisions, Decisions	173
<i>Glossary</i>	191
<i>Index</i>	197

Online Models

A model of infrared radiation in the atmosphere:

http://understandingtheforecast.org/Projects/infrared_spectrum.html

A model of visible + infrared radiation in the atmosphere:

http://understandingtheforecast.org/Projects/full_spectrum.html

A model of the geological carbon cycle:

<http://understandingtheforecast.org/Projects/geocarb.html>

Integrated assessment model (ISAM) for future climate change:

<http://understandingtheforecast.org/Projects/isam.html>

A Hubbert's Peak calculator:

<http://understandingtheforecast.org/Projects/hubbert.html>

The Kaya Identity model for the growth of the human footprint:

<http://understandingtheforecast.org/Projects/kaya.html>

Browsing the results of a coupled climate model:

<http://understandingtheforecast.org/Projects/bala.html>

The effect of Earth's orbit on sunlight intensity:

<http://understandingtheforecast.org/Projects/orbits.doc.html>

Preface

Global Warming: Understanding the Forecast is based on a class for undergraduate non-science majors at the University of Chicago, developed by Ray Pierrehumbert and myself. The class serves as partial fulfillment of our general education or core science requirements. We teach the class, and I have written the textbook, in a mechanistic way. We are aiming to achieve an intuitive understanding of the “ropes and pulleys” of the natural world, a fundamental scientific foundation that will serve the student for longer than would a straight presentation of the latest predictions.

The text is aiming at a single problem, that of assessing the risk of anthropogenic climate change. The story ranges from science to economics to policy, through physics, chemistry, biology, geology, computer science, and of course Earth sciences. We see the distant past and the distant future. In my opinion, by looking at one problem from many angles, the student gets a decent view of how a working scientist really thinks. This is as opposed to, say, taking a survey tour of some scientific discipline.

The text is suitable for students of all backgrounds. We do make some use of algebra, mostly in the form of what are known as (gasp!) story problems. The student will be exposed to science from many different fields, but no prior knowledge of any of these topics is required. One can learn something of what each field is about by learning what it is good for, within the context of the common unifying problem of global warming.

Each chapter after the first comes with exercises, either based on a computer model with a web interface a paper-and-pencil exercise, suitable to do in lab sections or as homework. The first three are paper and pencil, aimed at building a foundation for understanding the computer labs that follow. The models run on our computers at the University of Chicago, and you can access them through Web pages. No special setup of the student’s computer is required, and students can work equally well in a computer lab or at Starbuck’s (actually, it would be interesting to see if that’s true.).

This book benefited by thorough, thoughtful reviews by Andy Ridgwell, Stefan Rahmstorf, Gavin Schmidt, and a fourth, anonymous reviewer. The Web interface to the models benefited from input by Jeremy Archer. The visible-infrared radiation model was constructed by Ray Pierrehumbert and Rodrigo Caballero. The Integrated Science Assessment Model (ISAM) was provided by Atul Jain. The Exercises in Chapters 11 and 12 make use of model output provided by G. Bala and is plotted using Ferret, developed at the National Oceanic and Atmospheric Administration Pacific Marine Environmental Laboratory.

Preface to the Second Edition

The text has been revised and updated throughout, reflecting results from the Fourth Intergovernmental Panel on Climate Change Scientific Assessment Report, published in 2007.

**Everyone always complains about the weather,
but no one ever does anything about it**

—Mark Twain

The Glaciers Are Melting, But Is It Us?

Is it really possible that human activity could alter the weather? As I write, it is a crisp, clear fall day. What would be different about this day in a hundred years, if the climate of the Earth gets altered by human industrial activity?

There is no doubt that the Earth is warming. Mountain glaciers are disappearing. The Arctic coast is melting. Global average temperature records are broken year after year. The growing season has been getting longer. Plans are being made to abandon whole tropical islands as they sink into the Pacific Ocean. Shippers are awaiting the opening of the Northwest Passage that early explorers searched for in vain, with the melting of sea ice in the Arctic.

Of course, the natural world has variable weather all by itself, naturally. Is it likely that some of our recent weather has been impacted by human-induced climate change, or how much of this would have happened anyway? If humans are changing the climate, do we know that this is a bad thing? How does the future evolution of climate compare with the climate changes we may be seeing today, and natural climate changes in the prehistoric past?

Weather Versus Climate

We should distinguish at the outset between *climate* and *weather*. *Weather* is chaotic, which means that it cannot be forecast very far into the future. Small errors in the forecast grow with time, until eventually the forecast is nothing but error. The word *climate* means some kind of average of the weather, say over 10 years or more. Weather models cannot reliably predict whether it will rain or be sunny on a particular day very far into the future, but climate models can hope to forecast the average rainfall of some location at some time of year. Weather is chaotic, but the average is not chaotic and seems to be in some ways predictable (Chapter 6).

Human-induced climate changes are expected to be small compared with the variability of the weather. Temperature in the coming century is projected to rise by a few degrees centigrade (Chapter 12). This change is small compared with the temperature differences between the equator and the pole, between winter and summer, or even between daytime and nighttime. One issue this raises is that it is tricky to discern a change in the average, when the variability is so much greater than the trend. Careers are spent computing the global average temperature trends from the 100+-year thermometer record (Chapter 11). Another issue is that it makes it more difficult to know whether temperature trends are natural or human-driven.

The small change in the average, relative to the huge variability, also raises the question of whether a change in the average will even be noticeable. One way the average weather is

important is in precipitation. Groundwater tends to accumulate, reflecting rainfall over the past weeks and months. It might not matter to a farmer whether it rains on one day versus the next, but if the average rainfall in a region changes, that could spell the difference between productive farming and not. A change in the average climate would change the growing season, the frequency of extreme hot events, the distribution of snow and ice, the optimum growth localities of plants and agriculture, and the intensity of storms.

In addition to day-to-day weather, there are longer-lasting variations in climate. One past climate regime was the Little Ice Age, around 1650–1800, bringing variable weather to Europe. By some reconstructed temperature records, it was about 1°C colder than our “natural” climate from about the year 1950. Before that was the Medieval Climate Anomaly, perhaps 0.5°C warmer over Europe, coincident with periods of prolonged drought in the American Southwest. The causes of these climate changes are discussed in Chapter 11, but for now it is enough to observe that relatively small-sounding average-temperature shifts produced noticeable changes in human welfare and the evolution of history. The climate of the Last Glacial Maximum, 20,000 years ago, was so different from today that the difference would be obvious even from space, in the massive ice sheets and altered coastlines; and yet the average temperature difference between then and today was only about 5–6°C (Chapter 8).

Forecasting Climate Change

The fundamental process that determines the temperature of the Earth is the balance between energy flowing to the Earth from the sun versus energy flowing away from the Earth into space. Heat loss from Earth to space depends on Earth’s temperature (Chapter 2). A hotter Earth loses heat faster than a cooler one, everything else being equal. The Earth balances its energy budget by warming up or cooling down, finding the temperature at which the energy fluxes balance, with outflow equaling inflow.

It is possible to change the average temperature of the Earth by altering the energy budget, for example, by changing the brightness of the sun. The sun’s intensity covaries with the number of sunspots. Sometimes sunspots disappear altogether, presumably indicating a particularly cool sun. The Maunder Minimum was such a period, lasting from 1645–1750, coincident with the Little Ice Age.

Some of the incoming sunlight is reflected back to space without ever being absorbed (Chapter 7). Reflected sunlight leaves no heat energy behind, allowing the Earth to cool. Clouds reflect light, and so does snow. Bare soil in the desert reflects more light than vegetation does. Smoke emitted from coal-burning power plants produces a haze of sulfuric acid droplets that can reflect light.

Earth’s temperature is also sensitive to the outgoing energy flowing to space as infrared radiation. The greenhouse effect keeps the Earth a lot warmer than it would be without it. Most of the gases in the air are completely transparent to infrared light, meaning that they are not greenhouse gases. The greenhouse effect is driven entirely by trace gases in the atmosphere, primarily carbon dioxide (CO₂), water vapor, and methane (CH₄). The impact that a particular greenhouse gas has on climate depends on its concentration (Chapter 4). The strength of the greenhouse effect also depends on the temperature structure of the atmosphere (Chapter 5).

Water vapor is a tricky greenhouse gas because the amount of water vapor in the atmosphere is determined by the climate. Water tends to evaporate when the air is warm and condense as rain or snow in cool air. Water vapor, it turns out, amplifies the warming effects from changes in other greenhouse gases. This water-vapor feedback more or less doubles the temperature change we would expect from rising atmospheric CO₂ concentration without the feedback, in a dry world, for example.

Clouds are very effective at absorbing and emitting infrared light, acting like a completely infrared-opaque greenhouse gas. A change in cloudiness also affects the visible-light incoming energy flux by reflecting it (Chapter 7). Clouds are double-edged climate-forcing agents, cooling the Earth by reflecting sunlight while warming it by acting like a greenhouse gas in the atmosphere, trapping infrared light escaping to space. Changes in cloudiness are expected to amplify any climate warming in the future, but the strengths of this positive feedback are very uncertain. Clouds are the major uncertainty in the model calculation of the climate sensitivity of the Earth.

Human activity has the potential to alter climate in several ways. Rising CO₂ concentration from combustion of fossil fuel is the largest and longest-lasting human-caused climate-forcing agent, but we also release or produce other greenhouse gases, such as methane and other carbon molecules, nitrous oxide, and ozone. Particles and droplets arising from smokestacks and internal combustion engines reflect incoming visible light, altering the heat balance. Adding these into otherwise clean air may also change the average size of cloud droplets, which has a potentially huge but very uncertain impact on sunlight reflection (Chapter 10).

Many of these climate drivers themselves respond to climate, leading to stabilizing or destabilizing feedbacks. Reconstructions of prehistoric climate changes often show more variability than models tend to predict, presumably because there were positive feedbacks in the real world that are missing in the models. For example, the climate cools, so forest changes to tundra, allowing more of the incoming sunlight to be reflected to space, thereby cooling the climate even more. A climate model in which the forests do not respond to climate would underestimate the total amount of cooling. In the global warming forecast, the feedbacks are everything (Chapter 7).

The forecast for the coming century is also tricky because some parts of the climate system take a long time to change, such as melting an ice sheet or warming the deep ocean. It is hard enough to predict the equilibrium climate response to some change in forcing, but it is even harder to predict how quickly it will change (Chapter 12).

Carbon, Energy, and Climate

Climate change from fossil fuel combustion is arguably the most challenging environmental issue humankind has ever faced because CO₂ emission is at the heart of how we produce energy, which is pretty much at the heart of our modern standard of living. The agricultural revolution, which supports a human population of six billion people and hopefully more, has at its heart the industrial production of fertilizers, a very energy-intensive process. It is not so easy to stop emitting CO₂, and countries and companies that emit lots of CO₂ have a strong interest in continuing to do so (Chapter 9).

The energy we extract from fossil fuels originated in the nuclear fires of the sun. Visible light carried the energy to Earth, where it powered photosynthesis in plants, storing energy in chemical bonds between atoms of carbon, hydrogen, oxygen, and other elements. Plants have two motives for doing this: one to store energy and the other to build the biochemical machinery of life (Chapter 8).

Most of the biological carbon we use for fossil fuels was photosynthesized millions of years ago. Over geologic time, some of the biological carbon has been converted into the familiar fossil fuel types: oil, natural gas, and coal. Coal is the most abundant of these, but traditional oil will be depleted in a few decades (Chapter 9). Stored carbon energy is used to do work, in plants, animals, and now in automobiles, etc. by reacting the carbon with oxygen to produce CO₂. An automobile needs gas and air to run the engine, liberating the chemical energy by combustion. In living things the energy-extracting process is called *respiration*, and it explains why we need to breathe (to obtain oxygen and get rid of CO₂) and eat (to get biological carbon compounds) (Chapter 8).

CO₂ is released into the atmosphere to join the beautiful cacophony that is the carbon cycle of the biosphere. Trees and soils take up and release carbon, as does the ocean. Carbon is

released when tropical forests are cut, while forests in the high latitudes appear to be taking up atmospheric CO₂. Most of the CO₂ we release to the atmosphere will eventually dissolve in the oceans, but this process takes several centuries. A small fraction, about 10%, of the CO₂ released will continue to alter climate for hundreds of thousands of years into the future (Chapter 10).

Assessing the Risk

Is mankind creating a global warming trend? Climate scientists have tried to answer this question by comparing the history of Earth's temperature with the history of the different reasons why temperature might have changed, what are called *climate forcings*. The sun is more intense at some times than at others. Volcanoes occasionally blow dust or haze into the stratosphere, where it reflects sunlight back to space. Greenhouse gases and smokestack aerosols are two anthropogenic climate forcings.

Global temperature varies naturally, driven by changes in solar intensity and volcanic activity, but the Earth has been warming since the 1970s to an extent that can only be explained as global warming from greenhouse gases (Chapter 11). Greenhouse theory is more than a hundred years old and is required to explain why the natural Earth is as warm as it is (as well as other planets, such as Venus and Mars). The best guess for the climate sensitivity of the Earth, based on the thermometer record and on models, is also consistent with prehistoric climate changes documented in the geologic record. There is no existing model or theory of climate that can reproduce the present-day natural climate but would not warm significantly under business-as-usual CO₂ emission projections.

The forecast for the climate of the coming century is for a temperature increase of 2–5°C by the year 2100. It does not sound like much until it is compared with the impacts of natural climate changes in the past, such as the Little Ice Age (–1°C) and the Medieval Climate Anomaly (+0.5°C). These climate shifts were noticeable pretty much everywhere, and they were game-changing in some places.

The largest unknown in the climate forecast is the amount of CO₂ that will ultimately be released. Some amount of global warming is inevitable, but most of the carbon that could be released by 2100 is still in the ground. The decision is being made now and in the next few decades. The economic projections are that reducing CO₂ emissions substantially might cost a few percent of global net economic production, or GNP. That certainly would be a lot of money if you were looking at it piled in a heap, but in an economy that is growing by a few percent per year, a cost of a few percent per year would set the trajectory of economic production back by only a year or two (Chapter 13).

Sometimes climate issues are framed as a balance between the costs of avoiding CO₂ emission versus the costs of living with climate change. That way of thinking seems flawed to me because it ignores the unfairness that the people who benefit from CO₂ emission, mostly those in the developed world, are not the same people as those who pay the bill, in the developing world and in the distant future. Ultimately, the question of climate change may be a matter of ethics and fairness as much as one of profit and loss.

Greenhouse gas emission to the atmosphere is an example of a situation called the Tragedy of the Commons. The benefits of fossil-fuel combustion go to the individual, whereas the costs of climate change are paid by everyone. In this situation there is a natural tendency for everyone to overexploit the common resource. The solution to this is to change the incentive system so that the common costs of the decision end up being paid by the person making the decision. Carbon taxes and cap-and-trade schemes are two proposed ways to do this.

International negotiations are ongoing, within the Framework Convention on Climate Change (FCCC) under the United Nations. The first major milestone of the FCCC was the

Kyoto Protocol, initially adopted in 1995. It succeeded in entering into force in 2005, but it has failed to decrease the rate of CO₂ emission. The ongoing negotiations may be having an impact on energy investment, however, by stimulating the growth of alternative energy technology, for example.

In the next few decades, the biggest changes in carbon emissions could come from conservation and efficiency. The technology exists to begin to cut CO₂ emissions. Looking farther ahead, one hundred years from now, new, larger-scale energy sources will have to be found. Neither windmills nor nuclear power plants such as those we have today could scale up to the amount of carbon-free energy that humankind will need a century from now.

With our growing technological and intellectual prowess, as well as our exploding population, we are slowly taking over the job of managing the biosphere. May we do it wisely!

Part I

The Greenhouse Effect

How light can carry heat energy through empty space

Heat

Heat is simply the bouncing-around energy of atoms. Atoms in gases and liquids fly around, faster if it is hot and slower if it is cold. Atoms with chemical bonds between them stretch, compress, and bend those bonds, vibrating more energetically at high temperatures.

Perhaps you knew this already, but have you ever wondered how you can feel such a thing? Atoms are tiny things. We cannot really feel individual atoms. But it doesn't take any state-of-the-art laboratory technology to tell how fast the iron atoms in your stove burner are bouncing around. All you have to do is touch it. Actually, another method may be occurring to you—that you could look at it; if it is hot, it will glow. It glows with blackbody radiation, which we shall return to later.

A thermometer is like an atomic speedometer.

You can feel the hot stove because the energetic “bounciness” of the stove atoms gets transferred to the atoms in your nerves in your fingers. The fast-moving atoms of the burner bounce off the atoms in your finger, and the fast ones slow down a bit; the slower ones bounce back with a little more energy than they started with. Biological systems have evolved to pay attention to this, which is why you can feel it, because too much energy in our atoms is a dangerous thing. Chemical bonds break when they are heated too much; that's what cooking is. Burning your finger by touching a hot electric stove burner is an example of heat **conduction**, the easiest form of heat flux to envision.

Light

A thermos bottle is designed to slow the flow of heat through its walls. You can put warm stuff in there to keep it warm or cool stuff in it to stay cool. Thermos bottles have two walls, one inside and one outside. Between the two walls is an insulator. A **vacuum**, a space with no air in it, is a really good insulator because it has no molecules or atoms of gas in it to carry heat energy between the inner and outer walls of the thermos. Of course, there will still be heat conduction along the walls. But think about a planet. There are no walls connecting a planet with anything. The space between the Earth and the sun is a pretty good vacuum. We know how warm it is in the sunshine, so we know that heat flows from the sun to the Earth. And yet separating the sun from the Earth is 150 million kilometers of vacuum. How can heat travel through this?

Electromagnetic waves carry energy through a vacuum.

The energy travels by means of **light**. Electrons and protons have a property called **electric charge**. What an electric charge is, fundamentally, no one can tell you, but electric charges interact with each other through space via a property of the vacuum called the **electric field**. A positive electric field attracts a negatively charged electron. This is how a TV tube hurls electrons toward the screen in a picture-tube television, in the old days, before electronic flat screens, by releasing an electron and then pushing it around with the electric field.

The electric field interacts with another property of the vacuum called the **magnetic field**. If the strength of the electric field at some location, measurable in volts, changes with time, if the voltage is oscillating up and down, for example, that will cause a change in the magnetic field, such as a compass would point to. This is how an electromagnet works, converting electrical field energy into magnetic field energy. Going the other direction, if the magnetic field changes, it can produce an electrical field. This is how an electrical generator works.

It turns out that the electric and magnetic fields in the vacuum fit together to form a closed, self-replicating cycle like the ringing of a bell or the propagation of a nonbreaking wave in water. An up-and-down oscillation in the electric field will cause a complementary oscillation in the magnetic field, which reinforces the electric field in turn. The disturbances created by the electric and magnetic fields propagates the pattern in the fields, a self-reproducing wave moving through space at a particular velocity, the **speed of light**, written as the letter *c*. The bundle of electric and magnetic waves can, in principle, hurl through the vacuum forever, carrying energy with it.

A light wave rings like a bell.

The ringing of the electromagnetic field in light can come in a range of different frequencies. **Frequencies**, of oscillators or of light waves, have units of cycles per second and are denoted by the Greek letter ν , pronounced “nu”. It turns out that all different frequencies of light travel at the same speed in a vacuum. Within some nonvacuum media, such as air, water, or glass, different frequencies of light might vary in their speeds a little, which is how a prism separates white light into its component colors. But in a vacuum, all light travels at the same speed. The speed of light in a vacuum is a fundamental constant of nature.

The constancy of the speed of light in a vacuum makes it possible to relate the frequency of the light to its **wavelength**, the distance between the crests of a wave. You can figure out what the relationship is between frequency and wavelength by thinking geometrically, imagining the wavy line in Figure 2-1 to be moving past us at some speed, *c*. If the crests are 1 cm apart and they are moving at 10 cm/second, then 10 crests would move past us every second. Alternatively, you can make use of units. Pay attention to units, and they will lead you to virtue, or at least to the right answer most of the time. Problem: Assemble the two things we know, ν and *c*, such that the units combine to be the same as the units of λ . Solution:

$$\lambda \left[\frac{\text{cm}}{\text{cycle}} \right] = \frac{c \left[\frac{\text{cm}}{\text{sec}} \right]}{\nu \left[\frac{\text{cycle}}{\text{sec}} \right]}$$

Don’t take my word for it. Check and make sure the units are the same on both sides of the equation.




Figure 2-1 The frequency and wavelength of light are related to each other by the speed of light, which in a vacuum is the same for all different types of light, $3.0 \cdot 10^{10}$ cm/s.

Scientists who discuss infrared (IR) light often use a third way of describing different colors, called the **wave number**, defined as the number of cycles per centimeter of length. This is simply a matter of convenience; when IR light is described by its wave number, it will have a nice, easy-to-remember number; for example, the CO₂ absorption of IR light that is responsible for most of its greenhouse effect comes at a wave number value of about 700 cycles/cm. We will use the letter *n* to designate the wave number. How can we construct the units of *n* from the building blocks of λ , ν , and *c*? By looking at the units, it looks like all we need is λ :

$$n \left[\frac{\text{cycles}}{\text{cm}} \right] = \frac{1}{\lambda \left[\frac{\text{cm}}{\text{cycle}} \right]}$$

Different frequencies of light all have the same fundamental nature; they are all waves of the same essential physics. Figure 2-2 shows the names assigned to different types of light, based on their frequencies. Of course, if you know the frequency of light, you know its wavelength and wave numbers also, so the figure has “mile markers” in these units, too.

Our eyes are sensitive to light in what we pragmatically call the visible range. Light of higher frequencies than this carries enough energy to break chemical bonds, so it can be dangerous to us as chemical creatures. **Ultraviolet light**, or **UV**, causes sunburn, and **x-rays** and **gamma radiation** can do deeper and more extensive chemical damage. At longer wavelengths than the visible, we find the **infrared** range, or **IR**. Objects at room temperature glow with IR light. Heat lamps at the skating rink warm people up by shining invisible IR light on them.

All that energy whizzing around space in the form of coupled electric and magnetic field waves would be of little interest to the energy balance of a planet if the waves did not give up or carry off energy. There are a number of different mechanisms by which light may interact with matter, but IR light interacts mostly with vibrations of the chemical bonds of a molecule, coupled by means of the electric field that they share (Figure 2-3).

Imagine matter as constructed with charged oscillators all over its surface, little charged weights sticking out on little springs that stretch and contract. Each little oscillator will have a frequency with which it will “ring.” Incoming energy in the form of light brings with it an electric field oscillating up and down; voltage goes plus, minus, plus, minus. If the frequency of the cycles of the electric field from the light is about the same as the frequency of the oscillations of the chemical bonds, the light can be absorbed. Its energy is transferred into **vibrational energy** of the matter.




Figure 2-2 The electromagnetic spectrum.




Figure 2-3 A charged oscillator interacting with light.

This mechanism of energy transfer is a two-way street. If energy can flow from the light to the oscillator, it will also be able to go the other way, from the oscillator to light. The vibrational energy of the oscillator is what we have been calling its temperature. Any matter that has a temperature above absolute zero (zero degrees on the Kelvin scale) will have energy in its oscillators that it may be able to use to create light. The two-way street character of this process is important enough that it is given the name **Kirchoff's law**. Try an experiment of singing a single note into a piano with dampers off. When you stop, you will hear the strings echo back the note you sang into it. This sound wave to string vibration energy transfer is a two-way street as well.

Blackbody Radiation

Where can one see electrical energy traveling the other way, from matter into light? One example is a red-hot electric burner, which shines light you can see. The light derives its energy from the vibrations or thermal energy of the matter. We normally do not think of it, but it turns out that your electric burner continues to shine even when the stove is at room temperature. The difference is that the room temperature stove emits light in frequencies that our eyes cannot see, down in the IR.

If a chunk of matter has oscillators that vibrate and can interact with light at all possible frequencies, it is called a **blackbody**. The light that is emitted by a blackbody is called **blackbody radiation**. Most solids and liquids at the surface of the Earth are pretty good blackbodies, but gases in the atmosphere are not blackbodies; they only interact with specific frequencies of light. They are like pianos with most of the strings missing.

A blackbody is like a musical instrument with all the notes.

Blackbody radiation is made up of a characteristic distribution of frequencies (colors) of infrared light. Figure 2-4 shows a plot with axes of the intensity of light in the y -direction and frequency in the x -direction. The units of intensity look a bit complicated; they are $\frac{\text{Watts}}{\text{m}^2 \cdot \text{wave number}}$. The unit on the top of the fraction is **Watts**, the same kind of Watts that describe hairdryers and audio amplifiers. A Watt is a rate of energy flow, defined as Joules per second, where a Joule is an amount of energy, such as might be carried in a battery or a candy bar. The meters squared (m^2) on the bottom of the fraction is the surface area of the object. The unit of wave numbers on the bottom of the fraction allows us to divide the energy up according to the different wave-number bands of light; for instance, all the light between 100 and 101 cm^{-1} carries this many W/m^2 of energy flux, between 101 and 102 cm^{-1} carries that many W/m^2 , and so on. The total flux of energy can be calculated by adding up the bits from all the different slices of the light spectrum. The total energy of the wave number slice is I times the width of the slice in wave numbers, resulting in an energy flux in units of W/m^2 . You could cut the plot out with a pair of scissors and weigh the inside piece to determine its area, which would then be proportional to the total energy emitted in all frequencies of light.

A plot of intensity versus wavelength of light is called a spectrum. The IR light emission spectrum of a blackbody depends only on the temperature of the object. There are two things you should notice about the shapes of the curves in Figure 2-4. First, as the temperature goes up, the peaks of the curves move to the right, toward visible light. Second, as the temperature of the object goes up, the total energy emitted by the object goes up, which you can see by the fact that the areas under the curves in Figure 2-4 get larger as the temperature rises. There is an equation that tells how quickly energy is radiated from an object. It is called the **Stefan-Boltzmann equation**, and we are going to make extensive use of it. Get to know it now! The equation is:

$$I = \epsilon \sigma T^4 \quad (2-1)$$




Figure 2-4 The intensity of light emitted from a blackbody as a function of the wave number (cycles per centimeter) of the light. There are several spectra shown for different blackbody objects at different temperatures. A warmer object emits more radiation than a cooler one.

The **intensity** of the light is denoted by I and represents the total rate of energy emission from the object at all frequencies, in units of Watts/m². The Greek letter epsilon (ϵ) is the **emissivity**, a number between zero and one describing how good a blackbody the object is. For a perfect blackbody, $\epsilon = 1$, and the lower bound is $\epsilon = 0$. Sigma (σ) is a fundamental constant of physics that never changes, a number you can look up in reference books, called the **Stefan-Boltzmann constant**. T is the temperature in **Kelvins**, and the superscript 4 is an exponent, indicating that we have to raise the temperature to the fourth power. The Kelvin temperature scale begins with 0 K when the atoms are vibrating as little as possible, a temperature called **absolute zero**. There are no negative temperatures on the Kelvin scale.

A hot object emits much more light than a cold object.

One of the many techniques of thinking scientifically is to pay attention to units. Here is Equation 2-1 again, with units of the various terms specified in the square brackets:

$$I \left[\frac{\text{W}}{\text{m}^2} \right] = \epsilon [\text{unitless}] \sigma \left[\frac{\text{W}}{\text{m}^2 \text{K}^4} \right] T [\text{K}]^4$$

The unit of the intensity I is watts of energy flow per square meter. The meters squared on the bottom of that fraction is the surface area of the object that is radiating. The area of the Earth, for example, is $5.14 \cdot 10^{14} \text{ m}^2$. Temperature is in Kelvins, and ϵ has no units; it is just a number ranging from 0 to 1; 0 for an object that emits nothing and 1 for a perfect blackbody.



Figure 2-5 A photograph taken in infrared light.

Here is the important point: The units on each side of this equation must be the same. On the right-hand side, K^4 cancels leaving only W/m^2 to balance the left-hand side. In general, if you are unsure how to relate one number to another, the first thing to do is to listen to the units. They will guide you to the right answer. We will see many more examples of units in our discussion, and you may rest assured I will not miss an opportunity to point them out.

IR-sensitive cameras allow us to see what the world looks like in IR light. The clear lenses in the glasses of the guys in Figure 2-5 are cooler than their skins and, therefore, darker in the IR. How much darker they will be can be estimated from Equation 1 to be

$$\frac{I_{\text{cheek}}}{I_{\text{glasses}}} = \frac{\epsilon_{\text{cheek}} \sigma T_{\text{skin}}^4}{\epsilon_{\text{glasses}} \sigma T_{\text{glasses}}^4}$$

The Stefan-Boltzmann constant σ is the same on both top and bottom; σ never changes. The emissivity ϵ might be different between the skin and the glasses, but let us assume they are the same. This leaves us with the ratio of the brightnesses of the skin and glasses equal to the ratio of temperatures to the fourth power, maybe $(285\text{ K}/278\text{ K})^4$, which is about 1.1. The skin shines 10% more brightly than the surface of the coat, and that is what the IR camera sees.

The other thing to notice about the effect of temperature on the blackbody spectra in Figure 2-4 is that the peaks shift to the right as the temperature increases. This is the direction of higher-frequency light. You already knew that a hotter object generates shorter wavelength light because you know about red hot, white hot. Which is hotter? White hot, of course; any kid on the playground knows that. A room temperature object (say 273 K) glows in the IR, where we cannot see it. An electric stove set on High (400–500 K) glows in shorter wavelength light, which begins to creep up into the visible part of the spectrum. The lowest energy part of the visible spectrum is red light. Get the object hotter, say, the temperature of the surface of the sun (5,000 K), and it will fill all wavelengths of the visible part of the spectrum with light. Figure 2-6 compares the spectra of the Earth and the sun. You can see that sunlight is visible while “Earth light” (more usually referred to as **terrestrial radiation**) is IR. Of course, the total energy flux from the sun is much higher than it is from Earth. Repeating the calculation we used for the IR photo, we can calculate that the ratio of the fluxes is $(5,000\text{ K}/273\text{ K})^4$, or about 10^5 . The two spectra in Figure 2-6 have been scaled by dividing each curve by the maximum value that the curve reaches, so that the top of each peak is at a value of one. If we had not done




Figure 2-6 The shapes of the blackbody spectra of Earth and the sun. The Earth spectrum has been expanded to reach the same peak intensity as the solar spectrum, so the two can be compared on the same plot. The point is that the sun shines in visible light, whereas the Earth shines in infrared light.

that, the area under the Earth spectrum would be 100,000 times smaller than the area under the sun spectrum, and you would need a microscope to see the Earth spectrum on the figure.

Red hot, white hot.

It is not a coincidence that the sun shines in what we refer to as visible light. Our eyes evolved to be sensitive to visible light. The IR light field would be much more complicated for an organism to measure and understand. For one thing, the eyeball, or whatever light sensor the organism has, will be shining IR light of its own. The organism measures light intensity by measuring how strongly the incoming light deposits energy into oscillators coupled to its nervous system. It must complicate matters if the oscillators are losing energy by radiating light of their own. IR telescopes must be cooled to make accurate IR intensity measurements. Snakes are able to sense IR light. Perhaps this is possible because their body temperatures are colder than those of their intended prey.

TAKE-HOME POINTS

- Light carries energy through the vacuum of space.
- If an object can absorb light, it can also emit light.
- An object that can emit all frequencies of light is called a blackbody, and it emits light energy at a rate equal to $\epsilon \sigma T^4$.

STUDY QUESTIONS

1. Following the units, find the formula to compute the frequency of light given its wavelength or its wave number from its frequency.
2. Draw a blackbody radiation spectrum for a hot object and a cold object. What two things differ between the two spectra?

3. Use the Stefan-Boltzmann equation to compare the energy fluxes of the hot object and the cold object as a function of their temperatures.
4. What would an emission spectrum look like for an object that is not a blackbody?
5. How does the Stefan-Boltzmann equation deal with an object that is not a blackbody?

FURTHER READING

Blackbody radiation was a clue that something was wrong with classical physics, leading to the development of quantum mechanics. Classical mechanics predicted that an object would radiate an infinite amount of energy, instead of the $\epsilon \sigma T^4$, as we observe it to be. The failure of classical mechanics is called the ultraviolet catastrophe, and you can read about it,

lucidly presented but at a rather high mathematical level, in the *Feynman Lectures on Physics*, Volume 1, Chapter 41. My favorite book about quantum weirdness, the philosophical implications of quantum mechanics, is *In Search of Schrodinger's Cat* by John Gribbon, but there are many others.

Exercises

1. A joule (J) is an amount of energy, and a watt (W) is a rate of using energy, defined as $1 \text{ W} = 1 \text{ J/s}$. How many Joules of energy are required to run a 100-W light bulb for one day? Burning coal yields about $30 \cdot 10^6 \text{ J}$ of energy per kilogram of coal burned. Assuming that the coal power plant is 30% efficient, how much coal has to be burned to light that light bulb for one day?
2. A gallon of gasoline carries with it about $1.3 \cdot 10^8 \text{ J}$ of energy. Given a price of \$3 per gallon, how many Joules can you get for a dollar?

Electricity goes for about \$0.05 per kilowatt hour. A kilowatt hour is just a weird way to write Joules because a watt is a joule per second, and a kilowatt hour is the number of Joules one would get from running 1000 W times one hour (3,600 seconds). In the form of electricity, how many Joules can you get for a dollar?

A standard cubic foot of natural gas carries with it about $1.1 \cdot 10^6 \text{ J}$ of energy. You can get about $5 \cdot 10^5$ British thermal units (BTUs) of gas for a dollar, and there are about 1,030 BTUs in a standard cubic foot. How many Joules of energy in the form of natural gas can you get for a dollar?

A ton of coal holds about $3.2 \cdot 10^{10} \text{ J}$ of energy and costs about \$40. How many Joules of energy in the form of coal can you get for a dollar?

Corn oil costs about \$0.10 per fluid ounce wholesale. A fluid ounce carries about 240 dietary calories (which a scientist would call kilocalories). A calorie is about 4.2 J. How many Joules of energy in the form of corn oil can you get for a dollar?

Rank these as energy sources, cheap to expensive. What is the range in prices?
3. This is one of those job-interview questions to see how creative you are, analogous to one I heard: "How many airplanes are over Chicago at any given time?" You need to make stuff up to get an estimate and demonstrate your management potential. The question is: What is the efficiency of energy production from growing corn?

Assume that sunlight deposits 250 W/m^2 of energy on a corn field, averaging over the day-night cycle. There are 4.186 J per calorie. How many calories of energy are deposited on a square meter of field over the growing season? Now guess how many ears of corn grow per square meter, and guess what the number of calories is that you get for eating an ear of corn. The word *calorie*, when you see it on a food label, actually means *kilocalories* (thousands of calories), so if you guess 100 food-label calories, you are actually guessing 100,000 true calories, or 100 kcal. Compare the sunlight energy with the corn energy to get the efficiency.

4. The Hoover Dam produces $2 \cdot 10^9 \text{ W}$ of electricity. It is composed of $7 \cdot 10^9 \text{ kg}$ of concrete. Concrete requires 1 MJ of energy to produce per kilogram. How much energy did it take to produce the dam? How long is the “energy payback time” for the dam?

The area of Lake Mead, formed by Hoover Dam, is 247 mi^2 . Assuming 250 W/m^2 of sunlight falls on Lake Mead, how much energy could you produce if instead of the lake you installed solar cells that were 12% efficient?

5. It takes approximately $2 \cdot 10^9 \text{ J}$ of energy to manufacture 1 m^2 of crystalline-silicon photovoltaic cell. (Actually, the number quoted was 600 kWh. Can you figure out how to convert kilowatt hours into Joules?) Assume that the solar cell is 12% efficient, and calculate how long it would take, given 250 W/m^2 of sunlight, for the solar cell to repay the energy it cost for its manufacture.
6. We are supposed to eat about 2,000 dietary calories per day. How many watts is this?
7. Infrared light has a wavelength of about $10 \mu\text{m}$. What is its wave number in cm^{-1} ?

Visible light has a wavelength of about $0.5 \mu\text{m}$. What is its frequency in Hz (cycles per second)?

FM radio operates at a frequency of about 40 kHz. What is its wavelength?

Our first climate model

The Bare-Rock Model

The word **model** is used in scientific discussion to mean a fairly wide variety of things. Sometimes the word is synonymous with “theory” or “idea,” such as the Standard Model of Particle Physics. For doctors, a “model system” might be a mouse that has some disease that resembles a disease that human patients get. They experiment on the mouse rather than on people.

In climate science, models are used in two different ways. One way is to make **forecasts**. For this purpose, a model should be as realistic as possible and should capture or include all the processes that might be relevant in nature. This is typically a mathematical model implemented on a computer, although there is a nifty physical model of San Francisco Bay you should check out if you are ever in Sausalito. Once such a model has been constructed, a climate scientist can perform “what-if” experiments on it that could never be done to the real world, to determine how sensitive the climate would be to changes in the brightness of the sun or properties of the atmosphere, for example.

The simple model that we are going to construct here is not intended for making predictions, but it is more like a toy, to play with and understand and learn from. The model will demonstrate how the greenhouse effect works by stripping away lots of other aspects of the real world, things that would certainly be important for predicting climate change in the next century or the weather next week, but make the climate system more complicated and therefore more difficult to understand.

The layer model is a toy, demonstrating an idea.

The model is called the **Layer Model**. Understanding the Layer Model will not equip us to make detailed forecasts of future climate, but one cannot understand the workings of the real climate system without first understanding the Layer Model. The first step is to construct the model with no atmosphere, so that the only layer is the ground, a **Bare-Rock Model**. The next step is to demonstrate the greenhouse effect by adding an atmospheric layer.

The temperature of the surface of the Earth is controlled by the energy flows, in from the sun and back out to space as infrared (IR) light. Both types of light are primarily from blackbody radiation, as discussed in the last chapter. The sun is very bright because its surface temperature is high and therefore the energy flux $I = \epsilon \sigma T^4$ is a large number. Sunlight strikes the Earth and deposits some of its energy into heat at the Earth’s surface. The Earth sheds heat by blackbody radiating in the IR, at much lower intensity and longer wavelength because the Earth is cooler than the Sun.

The Layer Model assumes that the amount of energy leaving the Earth as IR is exactly equal to the amount of energy coming into the planet from sunlight. The real world may be out of energy balance for a little while or in some region or season, but we are assuming that the

Layer Model is exactly in balance. The Layer Model is essentially an algebraic story problem, and it can be built up from the algebraic statement that

$$F_{\text{out}} = F_{\text{in}} \quad (3-1)$$

where F_{out} is the outgoing energy flux, and F_{in} is the flux of energy in.

Let us begin with incoming sunlight. The **intensity** of sunlight, I_{in} , at the top of the atmosphere is about 1350 W/m^2 . But not all of this energy gets converted to planetary heat. A significant fraction of it is **reflected** rather than absorbed. If you have ever seen Venus shining brightly in the twilight sky, you will have seen a planet reflecting visible light. Venus's brightness is not blackbody radiation; it is hot on Venus but not hot enough to shine white hot. When light is reflected, for the purposes of the energy budget in the Layer Model, it is as if the energy had never arrived on Earth at all.

Reflected sunlight does not add to a planet's heat energy budget.

The fraction of a planet's incoming visible light that is reflected back to space is called the planet's **albedo** and is given the symbol α (Greek letter alpha). Snow, ice, and clouds are very reflective and tend to increase a planet's albedo. The albedo of Venus is high, 70%, because of a thick layer of sulfuric acid clouds in the Venusian atmosphere. Mars's albedo is low, about 15%, because of a lack of clouds on that planet. Earth's albedo is about 30% or 0.3, and it depends on factors that might change with changing climate, like clouds and ice.

Incoming solar energy that is not reflected is assumed to be absorbed into vibrational energy of molecules of the Earth. The Earthly albedo, α , of 0.3 means that the intensity of sunlight actually absorbed by the Earth is

$$I_{\text{absorbed}} = 1350 \text{ W/m}^2 \cdot (1 - \alpha) \approx 1000 \text{ W/m}^2$$

To get the total incoming flux for the whole planet, in units of Watts instead of W/m^2 , we need to multiply the intensity by a factor of area

$$F_{\text{in}}[\text{Watts}] = I_{\text{sunlight}} \left[\frac{\text{Watts}}{\text{m}^2} \right] \cdot (1 - \alpha) \cdot A[\text{m}^2]$$

but what area would it be? Sun shines on half of the surface of the Earth at any one time, but the light is weak on some parts of the Earth, near dawn or dusk or in high latitudes, and much more intense directly head-on to the sun, near the equator at noon. The difference in intensity is caused by the angle of the incoming sunlight, not because the sunlight, measured head-on at the top of the atmosphere, is much different between low and high latitudes (Figure 3-1). The sunlit side of the Earth is half its surface area, but using the area of a hemisphere to multiply by would require taking into account all the solar zenith angles from all the parts of the hemisphere.

The shadow knows.

But there is a trick that makes the calculation much simpler. The amount of head-on sunlight captured by the Earth can be easily calculated from the size of Earth's shadow. Use the formula for the area of a circle, not the surface area of a hemisphere. The area of a circle is

$$A[\text{m}^2] = \pi r_{\text{earth}}^2$$




Figure 3-1 When sunlight hits the Earth, it all comes from the same direction. The Earth makes a circular shadow. Therefore, the Earth receives an influx of energy equal to the intensity of sunlight, multiplied by the area of the circle, $\pi \cdot r_{\text{earth}}^2$.

Putting these together, the total incoming flux of energy to a planet by solar radiation is

$$F_{\text{in}} = \pi r_{\text{earth}}^2 (1 - \alpha) I_{\text{sunlight}} \quad (3-2)$$

The rate at which the Earth loses energy by radiation to space is given by the Stefan-Boltzmann equation:

$$F_{\text{out}} = A \epsilon \sigma T_{\text{earth}}^4$$

In the real world, the daytime side of the Earth is warmer than the nighttime side, and it is warmer near the equator than near the poles. The temperature differences from place to place would be even more dramatic on an airless planet like the moon. But the Layer Model is going to ignore that for now and will find a single value for the temperature of the Earth to go along with a single value for each of the heat fluxes, F_{in} and F_{out} . The effects of all the heterogeneity of the Earth's surface are explored in Chapter 6.

As for solar energy, to convert intensity I , in W/m^2 , to the total energy flux in watts, we need to multiply by an area A . What area is appropriate this time? Earthlight is different from sunlight in that sunlight is all traveling in one direction, whereas Earthlight leaves Earth in all directions (Figure 3-2). Therefore, the area over which the Earth radiates energy to space is simply the area of the sphere, both sides, day and night. The surface area of a sphere is given by

$$A_{\text{sphere}} = 4\pi r^2$$

and so the total outgoing energy flux from a planet by blackbody radiation is

$$F_{\text{out}}[\text{Watts}] = 4\pi r_{\text{earth}}^2 \epsilon \sigma T_{\text{earth}}^4 \quad (3-3)$$

The surface of the Earth is exactly four times bigger than the size of its shadow!




Figure 3-2 When infrared light leaves the Earth, it does so in all directions. The total rate of heat loss equals the intensity of the Earthlight multiplied by the area of the surface of the sphere $4 \cdot \pi \cdot r_{\text{earth}}^2$.

Setting the energy fluxes F_{out} and F_{in} equal to each other, from Equations 3-2 and 3-3, we get

$$4\pi r_{\text{earth}}^2 \varepsilon \sigma T_{\text{earth}}^4 = \pi r_{\text{earth}}^2 (1 - \alpha) I_{\text{in}}$$

which is essentially an expanded version of Equation 3-1. There are a lot of terms in this equation, but it can be simplified. There are factors of π and r_{earth}^2 on both sides, which can be cancelled. Dividing also by 4 on both sides puts the equation into units of W/m^2 of the Earth's surface,

$$\varepsilon \sigma T_{\text{ground}}^4 = \frac{(1 - \alpha) I_{\text{in}}}{4} \quad (3-4)$$

It is as though full-strength sunlight only shines on exactly one-fourth of the Earth.

We know everything in this equation except T_{Earth} . Rearranging the equation to put what we know on the right-hand side, and what we do not know on the left, we get

$$T_{\text{ground}} = 4 \sqrt[4]{\frac{(1 - \alpha) I_{\text{in}}}{4 \varepsilon \sigma}} \quad (3-5)$$

where the number 4 by the root sign indicates a fourth root instead of the usual square root.


This equation ties together several climate-sensitive quantities. Changes in solar intensity such as the sunspot cycle or the Maunder Minimum (Chapter 10) may affect I_{in} . The albedo of the planet is very sensitive to ice and cloud cover, both of which might change with changing climate. And the next chapter shows that greenhouse gases are extremely selective about the wavelengths of light that they absorb and emit; in other words, they have complexities relating to their emissivity ε values.

Solving Equation 3-5 for the temperature of the Earth gives a value of 255 K, or about -15°C . This is too cold. The average temperature of the Earth is closer to $+15^\circ\text{C}$ than it is to -15°C . Table 3.1 gives the values needed to do the same calculation for Venus and Mars, along

Table 3.1 Temperatures and Albedos of the Terrestrial Planets.

	$I_{\text{solar}}, \text{ W/m}^2$	$\alpha(\%)$	$T_{\text{bare}} \text{ K}$	$T_{\text{observed}} \text{ K}$	$T_{\text{1 layer}} \text{ K}$
Venus	2,600	71	240	700	285
Earth	1,350	33	253	295	303
Mars	600	17	216	240	259

The intensity of sunlight differs with distance from the sun.

**Figure 3-3** An energy diagram for Earth with no atmosphere, just a bare rock in space.

with the results of the calculation and the observed average temperatures. In all three cases, the Bare-Rock Layer Model has gotten it too cold.

The Layer Model with Greenhouse Effect

The Bare-Rock Model is too cold because it lacks the **greenhouse effect**. Like the rest of the Layer Model, the atmosphere that we are going to add is absurdly simple. It is essentially a pane of glass suspended by magic above the ground (Figure 3-3). Like glass, our atmosphere is transparent to visible light, so the incoming energy from the sun passes right through the atmosphere and is deposited on the planet surface, as before. The planet radiates energy as IR light according to $\epsilon\sigma T_{\text{ground}}^4$, as before.

In the IR range of light, the pane of glass acts as a blackbody, absorbing and emitting all frequencies of IR light. Therefore, the energy flowing upward from the ground, which we will call $I_{\text{up, ground}}$, is entirely absorbed by the glass. The glass in turn radiates energy according to $\epsilon\sigma T_{\text{glass}}^4$. Because the pane of glass has two sides, top and bottom, it radiates energy both upward and downward, $I_{\text{up, glass}}$ and $I_{\text{down, glass}}$.

The glass is transparent to visible light but “black” in the infrared.

It is more complicated to solve the Layer Model when the pane of glass is included because there are two unknowns to solve for, the temperatures of the ground and of the glass. Both these




Figure 3-4 An energy diagram for a planet with a single pane of glass for an atmosphere. The glass is transparent to incoming visible light but a blackbody to infrared light.

layers are in a state of energy balance, where the rate of energy going out of the layer equals the rate of energy coming in:

$$I_{\text{out}} = I_{\text{in}} \quad (3-6)$$

There is only one temperature for each of the layers at which both energy budgets will be in balance.

The energy fluxes are depicted as arrows in Figure 3-4. The **energy budget for the pane of glass** can be written as

$$I_{\text{up, glass}} + I_{\text{down, glass}} = I_{\text{up, ground}}$$

Because these are both IR light energy, the intensities can be written as

$$2\varepsilon\sigma T_{\text{glass}}^4 = \varepsilon\sigma T_{\text{ground}}^4 \quad (3-7)$$

Both the ground and the glass have to be in a state of energy balance.

The **budget for the ground** is different from in the Bare Rock Model because now there is heat flowing down from the pane of glass. The basic balance is

$$I_{\text{up, ground}} = I_{\text{in, solar}} + I_{\text{down, glass}}$$

These intensities can be expanded into

$$\varepsilon\sigma T_{\text{ground}}^4 = \frac{(1-\alpha)}{4} I_{\text{solar}} + \varepsilon\sigma T_{\text{glass}}^4 \quad (3-8)$$

Algebraically, it is possible to combine Equations 3-7 and 3-8 in such a way as to eliminate one of the two temperatures, solving for the other, and then plugging back in to solve for the first. But there is a nifty trick that makes it algebraically much easier to solve and is conceptually useful to boot.

The trick is to construct an **energy budget for the Earth overall** by drawing a boundary above the atmosphere and figuring that if energy gets across this line coming in, it must also be flowing across the line out at the same rate. This budget can be written as

$$I_{\text{up, glass}} = I_{\text{in, solar}}$$

The intensities comprise individual fluxes from the sun and from the atmosphere:

$$\epsilon\sigma T_{\text{glass}}^4 = \frac{(1 - \alpha)}{4} I_{\text{solar}}$$

This equation contains only one unknown, T_{glass} . Notice also that this equation looks a lot like Equation 3-4, describing the surface temperature of the Bare-Rock model. Solving for T_{glass} here gives the same answer as solving for T_{ground} in the Bare Rock model.

The skin temperature of the model is always the same.

This is actually more than just a curiosity or an algebraic convenience. What it means is that the place in the Earth system where the temperature is the most directly controlled by incoming sunlight is at the top layer, where infrared radiates to space. We will call this temperature the **skin temperature** of the Earth, and it comes up again in the next chapter.

What about temperatures below the skin, in this case T_{ground} ? Now that we know that the outermost temperature, T_{glass} , is equal to the skin temperature, we can plug that into the budget equation for the atmosphere to see that

$$2\epsilon\sigma T_{\text{glass}}^4 = \epsilon\sigma T_{\text{ground}}^4$$

or

$$T_{\text{ground}} = \sqrt[4]{2} T_{\text{glass}}$$

The temperature of the ground must be warmer than the skin temperature, by a factor of the fourth root of two, an irrational number that equals about 1.189. The ground is warmer than the atmosphere by about 19%. The One-Layer Model calculation is also repeated for Venus and Mars in Table 3.1. It gets the temperature of Earth about right, whereas Mars is too warm and Venus is not yet warm enough.

The glass layer warms the ground by trapping outgoing heat, like a winter coat traps outgoing body heat.

The Layer Model was never intended to be all that realistic. It was intended to illustrate the mechanism of the greenhouse effect, making it easier to think about and understand. So let's think about how it works. The ground gets warmer with the pane of glass installed, but the pane

of glass does not bring any new energy to the Earth. It is not like a humungous heat lamp in the sky. How, then, does it change the temperature of the ground?

The situation can be understood by analogy to water flowing through a kitchen sink. Water flows into the sink from a faucet at some rate, and it flows down the drain at a rate which depends on how much water is in the sink. The kitchen sink analogy will come up multiple times in this book, applied to multiple situations. For now, the water is analogous to energy flowing into and out of the Earth system.

Water drains faster as the water level in the sink rises because there is higher pressure under deeper water to push the water down the drain faster. This is just like the outgoing IR light energy flux, which flows faster as the temperature of the Earth increases, according to $\epsilon\sigma T^4$. The water flowing into the sink represents sunlight, corrected for reflection (albedo), which we could imagine as some fraction of the water from the faucet splashing out of the sink, not going down the drain. Like sunlight in Earth's energy balance, the water from the faucet acts as a driver to the kitchen sink water cycle, coming in from outside at a rate that is set externally rather than responding to the sink in some way like the water flow down the drain does.

Eventually, the water in the sink reaches a level where the outflow of water balances the inflow. That's the equilibrium value, and it is analogous to the equilibrium temperature calculated by the Layer Model. For the sink, we could ask what the water level should be where the flow down the drain balances the inflow from the faucet.

The greenhouse effect is like a constriction in the drain, like a piece of carrot that becomes lodged in the drain filter. For a few minutes after the carrot lodges on the drain filter the water drains out more slowly than the inflow from the faucet, so the water level in the sink rises. Eventually the water level reaches the point where water goes down the drain fast enough to balance the water budget again, with higher water level compensating for the constriction. A greenhouse gas, like the carrot in the drain filter, makes it more difficult for the heat to escape the Earth. The temperature of the Earth rises until the fluxes balance again.

TAKE-HOME POINTS

- The outflow of IR energy from a planet must balance heating from the sun.
- The planet accomplishes this act of energetic housekeeping by adjusting its temperature.
- Absorption of outgoing IR light by the atmosphere warms the surface of the planet as the planet strives to balance its energy budget.

STUDY QUESTIONS

1. A greenhouse layer increases the amount of heat energy in the ground layer. Where does this energy come from?
2. Relate the pieces of the kitchen sink analogy, the water level and the flows through the faucet and the drain, to different parts of the energy cycle through the Layer Model.
3. What factors in the Layer Model could affect the skin temperature?
4. What factors can change the water level in the kitchen sink analogy?

FURTHER READING

Kump, L. R., J. F. Kasting, and R. G. Crane. *The Earth System*. Saddle River, NJ: Prentice Hall, 1999.

Kasting J. F., and D. Catling. "Evolution of a habitable planet". *Annual Review of Astronomy and Astrophysics*. 41:429–463, 2003.

Exercises

1. **The moon with no heat transport.** The Layer Model assumes that the temperature of the body in space is all the same. This is not really very accurate, as you know that it is colder at the poles than it is at the equator. For a bare rock with no atmosphere or ocean, like the moon, the situation is even worse because fluids like air and water are how heat is carried around on the planet. So let us make the other extreme assumption, that there is no heat transport on a bare rock like the moon. Assume for comparability that the albedo of this world is 0.33, same as Earth's. What is the equilibrium temperature of the surface of the moon, on the equator, at local noon, when the sun is directly overhead? What is the equilibrium temperature on the dark side of the moon?
2. **A Two-Layer Model.** Insert another atmospheric layer into the model, just like the first one. The layer is transparent to visible light but a blackbody for infrared.
 - a) Write the energy budgets for both atmospheric layers, for the ground, and for the Earth as a whole, like we did for the One-Layer Model.
 - b) Manipulate the budget for the Earth as a whole to obtain the temperature T_2 of the top atmospheric layer, labeled Atmospheric layer 2 in Figure 3-5. Does this part of the exercise seem familiar in any way? Does the term ring any bells?
 - c) Insert the value you found for T_2 into the energy budget for layer 2, and solve for the temperature of layer 1 in terms of layer 2. How much bigger is T_1 than T_2 ?
 - d) Now insert the value you found for T_1 into the budget for atmospheric layer 1 to obtain the temperature of the ground, T_{ground} . Is the greenhouse effect stronger or weaker because of the second layer?




Figure 3-5 An energy diagram for a planet with two panes of glass for an atmosphere.




Figure 3-6 An energy diagram for a planet with an opaque pane of glass for an atmosphere.

- 3. Nuclear Winter.** Let us go back to the One-Layer Model but change it so that the atmospheric layer absorbs visible light rather than allowing it to pass through (Figure 3-6). This could happen if the upper atmosphere were filled with dust. For simplicity, assume that the albedo of the Earth remains the same, even though in the real world it might change with a dusty atmosphere. What is the temperature of the ground in this case?

Why some gases are greenhouse gases, but most aren't, and some are stronger than others

About Gases

The layer model is what is called an **idealization** of the real world. It has the essential ingredient of the greenhouse effect, but it is missing numerous things that are important in the real atmosphere. Starting from the Layer Model, the next few chapters add things one at a time, assembling the rest of the ingredients that control the temperature of the real Earth. The first task is to understand how real gases interact with infrared (IR) light in the real atmosphere.

Let us begin by defining different ways of describing the amounts of gases in the atmosphere. The word **concentration** means the number of molecules within some volume. The difficulty this raises for gases in the atmosphere is that the concentration, defined in this way, changes as the gas expands or contracts. It is often more convenient to talk about proportions of gases; for example, oxygen is about 20% of the molecules of gas in the atmosphere, and nitrogen is almost 80%. The proportion of CO₂ is currently 0.039%. We can express that in a more mnemonic way by saying 390 parts per million, or **ppm**. This number is called a **mixing ratio**.

A gas exerts pressure, a pushing force, on itself or on the surface of anything in contact with it. The force comes from deflected momentum of gas molecules bouncing off the surface. The pressure of a mixture of gases can be taken apart into what are called **partial pressures** arising from each type of molecule in the gas, which add up to the total pressure. The partial pressure of CO₂, for example, is written as **pCO₂**.

One might expect heavier gases to contribute disproportionately to the pressure because they seem like they would hit the walls harder. It works out, however, that heavier gas molecules are moving more slowly than the light ones are, just enough so that the pushing force, proportional to the mass times the velocity, is the same on average with every type of molecule regardless of how heavy it is. It is as if a bunch of ping-pong balls and bowling balls rolled against the wall all pushed equally on the wall. They are equal because you have to roll the bowling balls more slowly. At a given temperature, each molecule has the same amount of energy on average invested in motion, what is called **kinetic energy**. A heavy molecule will move more slowly with its allotment of kinetic energy.

A gas pressure, in units of microatmospheres (μatm), is numerically nearly equal to its mixing ratio in parts per million (ppm).

The bottom line is that the partial pressure from a gas is more or less proportional to the mixing ratio, the numerical proportion of the gas. If you think in pressure units of atmospheres, the actual numbers are the same for the partial pressure and the mixing ratio. For CO₂, for example, the mixing ratio is currently about 390 ppm, and its pCO₂ is about 390 μatm .

Gases, Vibrations, and Light

Most of the mass of an atom is in its nucleus, which resembles conceptually a massive sun at the center of a solar system. Two nuclei of two different atoms always repel each other because of their positive charges. Electrons float in ghostly quantum mechanical probability clouds, called **orbitals**, around the nucleus. For some combinations of atoms, the electrons fall into orbitals that are shared between the two nuclei. These electrons are the glue that holds together a **chemical bond** between the atoms.

A chemical bond is like a spring in that the two nuclei on either end of the bond have some freedom to move closer or farther apart. There is an optimum, lowest-energy distance for the nuclei to be from each other. If they are any closer, the positive nuclei will start to repel each other. Farther from each other, you get less energy gain from sharing the electrons. A bond vibrates when the distance between the nuclei oscillates between the nuclei being too close together, then too far apart. This is called a **stretching vibration**.

An atom may participate in more than one chemical bond. The bonds, and extra electrons that the central atom holds on to, arrange themselves around the atom to find their positions and angles of least energy. Here again there is some flexibility in the angle of the bond, which leads to a **bending vibration**.

Chemical bonds vibrate at particular frequencies.

Molecular vibrations in gas molecules are like vibrations of a piano string in that they tend to be at specific “ringing” frequencies. For weights and springs, the vibrational frequencies depend on two things: the mass of the weights and the springiness of the spring holding them together. Imagine the weights moving together, nearing the optimal distance but with enough momentum to carry them closer than they would like to be, continuing the oscillation. Heavier weights would have enough inertia to keep moving closer for a longer time than would a lighter pair of weights, so the frequency of the vibration will be slower with heavy weights. If the spring is very stiff, it will reverse the velocity of a vibration more quickly, and the frequency of the oscillation will be higher. Vibrations in chemical bonds depend on the mass of the nuclei and on the energy penalty for stretching or bending the bond, which determines the springiness of the chemical bond.

Most forms of liquid and solid molecules are not so “choosy” about frequency of IR light that they can absorb or emit, because the molecules interact with each other so much that more or less all frequencies of IR light can be absorbed and emitted. Liquids and solids are pretty good blackbodies, in other words. Gases however only absorb and emit specific frequencies of IR light, meaning that they are generally not very good blackbodies.

In fact, most of the gases in the atmosphere do not absorb or emit IR light at all, because vibrations in their bonds do not create an imbalance in the electrical field. Both O_2 and N_2 , the most abundant gases in the atmosphere, are symmetric molecules, made of two identical atoms whose electric fields just cancel each other out. Neither atom can hold the electrons any more tightly than the other, so there is no difference between the electrical field on one side of the molecule versus the other. The symmetry is unbroken when you stretch or compress the bond. Symmetrical molecules with only two atoms are never greenhouse gases.

To create or absorb infrared light, the molecule must be electrically lopsided, at least in passing.

If the symmetry is broken by having different types of atoms on each side, like carbon monoxide (CO) or nitrogen oxide (NO), the molecules begin to have some greenhouse properties. Both these types of molecules are very reactive and are not found in enough abundance in the atmosphere to lead to much greenhouse effect, but these are good stepping-stone molecules to think about before moving to more complicated atoms like CO₂. Because the atoms have different charges in their nuclei, their chemistry is different, including how hard they pull on those electrons in their shared chemical bond. One side of the molecule will probably have a slight positive charge, and the other will be slightly negative. This charge imbalance is called a **dipole moment**. When the bond vibrates, it changes the steepness of the transition from positive to negative, which drives the electric field. Changing the electric field is how greenhouse gases interact with IR light. One could even oscillate the electric field by simply rotating the molecule.

Molecules with three or more atoms must have more than one chemical bond because a single bond never connects more or less than two atoms. Imagine a set of three weights connected in a line with two springs. When this contraption is jigged, its bonds bend and stretch together, as an integrated system, rather than each bond simply vibrating according to its own springiness and weight, independent of the other. The whole thing rings like a bell, and like a complex bell with a rich tone, there are multiple oscillations with different frequencies all going on at once.

The CO₂ molecule is shaped in a straight line with carbon in the middle (Figure 4-1). Oxygen tends to pull on electrons more tightly than carbon does, but the oxygen atom on one side pulls the electrons just as tightly as the other oxygen on the other side. The molecule has no permanent electrical field asymmetry (dipole moment). If all you do to the molecule is rotate or stretch it, you have not broken the symmetry and you will have no effect on the electric field.

The CO₂ bending vibration mode is the important one for climate.

However, two modes of vibration do generate an asymmetry in the electric field. The most important CO₂ vibrational mode for climate is the **bending vibration**. When the CO₂ molecule is bent, the oxygens, carrying slight negative charges, swing from one side of the molecule to the other. The CO₂ bending vibration absorbs and emits IR light; it is said to be **infrared active**. The next section shows that the bending vibration absorbs IR light in the heart of Earth's emission spectrum, giving it strong leverage over Earth's climate.




Figure 4-1 Vibrational modes of a CO₂ molecule that interact with infrared light in the atmosphere.




Figure 4-2 Vibrational modes of a water molecule that interact with infrared light in the atmosphere.

The other vibrational mode is an **asymmetric stretch**, in which one bond is growing longer as the other gets shorter, back and forth. There is less IR light of this frequency to absorb in the atmosphere, so this mode has less impact on the Earth's radiation budget than the bend does.

Water, H_2O , is a molecule that is bent in its lowest energy state (Figure 4-2). This is because the oxygen has two pairs of electrons hanging off it, which push the hydrogens toward the other side. Hydrogen atoms hold their electrons more loosely than oxygen atoms in chemical bonds, so each hydrogen has a slightly positive charge (denoted in Figure 4-2 as the lowercase Greek letter delta $\delta+$). The oxygen end of the molecule has a slight negative charge. In contrast to CO_2 , water has a dipole moment built into its resting structure. Just as for the NO molecule, rotating an H_2O molecule would oscillate the electric field and generate light. Because the arrangement of the nuclei in H_2O is more complex than for NO, there are many modes of vibration of the water molecule, including a symmetric stretch and a bend. These modes are also IR active.

Water vapor is very electrically lopsided and can absorb and emit lots of frequencies of infrared light.

How a Greenhouse Gas Interacts with Earthlight

Figure 4-3 shows blackbody spectra for temperatures ranging from 300 K, surface temperature on a hot summer day, down to 220 K, which is about the coldest it gets in the atmosphere, up near the troposphere at about 10-km altitude. There is also a jagged-looking curve moving among the smooth ones. This is the intensity of light that an IR spectrometer would see if it were in orbit over the Earth, looking down. Figure 4-3 is not data, but rather a model simulation from one of our online models that will be used in the Exercises. The Web address is http://understandingtheforecast.org/projects/infrared_spectrum.html.

The spectrum of the light leaving the Earth going into space ranges between two different blackbody spectra, a warmer one of about 270 K and a colder one from about 220 K. Parts of the spectrum, most notably the broad, smooth part around 1,000 cycles/cm, follow the warmer blackbody spectrum. There do not happen to be any gases in the atmosphere that absorb or emit at this frequency, so this IR light is coming directly from the ground. Because the greenhouse gases in the atmosphere are transparent to IR light in this band of frequencies, it is called the **atmospheric window**.

The blackbody curves are like the temperature scale on a thermometer. You can read the temperature of the infrared source as its intensity against this scale.




Figure 4-3 The solid line is a model-generated spectrum of the infrared light escaping to space at the top of the atmosphere. For comparison, the broken lines are blackbody spectra at different temperatures. If the Earth had no atmosphere, the outgoing spectrum would look like a blackbody spectrum for 270 K, between the 260 K and 280 K spectra shown. The atmospheric window is between about $900\text{--}1,000\text{ cm}^{-1}$, where no gases absorb or emit infrared light. CO_2 , water vapor, ozone, and methane absorb infrared light emitted from the ground and emit lower-intensity infrared from high altitudes, where the air is colder than at the surface.

Other parts of the spectrum seem to follow the colder blackbody curve, like, for example, at around 700 cycles/cm. This IR light is coming from greenhouse gases in the cold upper atmosphere. The big dip in the spectrum at about 700 cycles/cm is from the bending vibration of CO_2 . IR light emission and absorption are a two-way street (Chapter 1), so CO_2 absorbs the intense, “warm” IR from the ground at this frequency and re-emits the same frequency but at a lower intensity, because of its colder temperature (Figure 4-4). The atmosphere is called **optically thick** in the CO_2 bend frequency range and **optically thin** in the atmospheric window.

The situation is analogous to standing on a pier and looking down into a pond of water. If the water were very clear, you could see light coming from the bottom; you would see rocks or old tires in the reflected light. If the water were murky, the light you would see would be scattered light coming from perhaps just a few inches down into the water. The old tires would be invisible, alas. The muddy lake is optically thick, so the light you see comes from near the top.

Figure 4-3 is constructed so that the area under a curve of a spectrum is proportional to the total energy flux. You can eyeball the energy change as the area change. The areas of the pure blackbody curves go up proportionally to the temperature raised to the fourth power because of the Stefan-Boltzmann equation (Equation 2-1 in Chapter 2). The area trick works with the jagged spectrum as well, representing the total energy loss of the planet to space. The effect of an atmospheric absorption band is to take a bite out of the blackbody spectrum from the Earth’s surface, decreasing the area and therefore decreasing the outgoing energy flux.




Figure 4-4 A comparison of the fate of infrared light in the optically thick CO₂ bend frequency (*left*) versus the optically thin atmospheric window (*right*).

Methane absorbs IR light at around 1300 cycles/cm. Whereas the CO₂ bending vibration absorbs right at the peak of the blackbody spectrum of the Earth, methane absorbs in one of the tails of the spectrum. In a competition between CO₂ and methane for which is the stronger greenhouse gas, the advantage here goes to CO₂. However, the next section introduces another difference between the gases, which more than makes up for this, so that methane winds up as the more powerful greenhouse gas of the two.

The Band Saturation Effect

Figure 4-5 shows the results from a series of model runs for different atmospheric CO₂ concentrations, ranging from 0 to 1000 ppm. For reference, the atmosphere in 2010 contains about 390 ppm. With no CO₂, the atmosphere would be transparent to light of around 700 cycles/cm, like it is in the atmospheric window. Adding the first 10 ppm of CO₂ has a fairly noticeable impact on the shape of the outgoing light spectrum, a sharp V-shaped slice out of the spectrum. As the CO₂ concentration is increased to 100 ppm, the center of the peak runs into the blackbody curve from just a bit colder than 220 K, and it does not get any deeper as the CO₂ concentration is raised to 1000 ppm. This is the **band saturation effect**, the band referring to the range of frequencies, and saturation meaning “used up,” all absorbed by the CO₂.

The reason for the saturation effect can be understood by analogy to the murky pond. As the water gets murkier, it decreases the distance that a photon of light can travel before it is absorbed. It does not take much mud in the water to obscure the old tire on the bottom, shifting the visible depth from the bottom at say 3 m to maybe only 1 m. If the pond gets a lot murkier, you will be able to see only a few centimeters down into the water. Making it murkier still will limit the view to only 1 cm. The change in depth is getting less sensitive to the murkiness of




Figure 4-5 A demonstration of band saturation by CO₂. The addition of 10 ppm CO₂ (*upper right*) makes a huge difference to the outgoing infrared light spectrum relative to an atmosphere that has no CO₂ (*upper left*). Increasing CO₂ to 100 and 1,000 ppm (*lower panels*) continues to affect the spectrum, but you get less bang for your CO₂ buck as CO₂ concentration gets higher.

the pond. In the same way, the changes in the temperature at which the atmosphere radiates to space get smaller as the CO₂ concentration of the air gets higher. You just see the coldest light that you can get.

The center of an absorption band gets saturated as the gas gets more concentrated, but the band gets fatter.

Each model run returns a value for the total IR energy flux leaving the location, called I_{out} on the webpage and in units of W/m². These values are plotted in Figure 4-6. There is a huge change in the outgoing energy flux between the runs with 0 ppm and 10 ppm because you go from having no absorption to having a quite noticeable peak by adding a relatively small amount of CO₂, just 10 ppm. Then, as you go up to 100 and 1,000 ppm, the change in energy balance for each new molecule of CO₂, or each new ppm of CO₂, gets weaker.

The energy balance never becomes totally insensitive to further additions of CO₂, however. More CO₂ will always make it warmer, even up to the 70 atm of CO₂ in the atmosphere of Venus. In part, this is because as the CO₂ concentration increases, the absorption peak in Figure 4-5 gets fatter, even if it cannot get much deeper. The middle of the absorption band is saturated, but the edges are unsaturated.




Figure 4-6 Band saturation viewed in a different way from Figure 4-5. This is a plot of the total energy flux carried by all infrared light, which is proportional to the area under the spectrum curves in Figure 4-5. The outgoing energy flux is less sensitive to CO₂ when the CO₂ concentration is high.

Even though the center of the CO₂ bend frequency is saturated, more CO₂ would still affect Earth's infrared energy output.

Given that the absorption bands are determined by vibrational frequencies of the molecule, one might have expected that the edges of the bands would be quite sharp, as if CO₂ absorbed 600 cycles/cm light completely and 599 cycles/cm light not at all. There are several reasons for the fuzziness in the peaks. One is **Doppler shifting**. Like a train whistle, which sounds a higher pitch as the train approaches, then lower after it passes and is moving away, the frequency of light that the molecule “sees” depends on the speed of the molecule, toward or away from the IR light source. If the molecule is moving away from the light source, it will absorb a slightly longer wavelength of light.

The absorption bands are also “smeared out” when gas molecules interact with each other, a phenomenon called **pressure broadening**. Liquids and solids are often fairly good blackbodies because the molecules are so close to each other that they share energy. Even if the frequency of the light is different from the frequency of the vibration, if the molecule is bathed in other molecules, it may be possible to take up the “slop” and absorb the light anyway. Gas molecules are generally too far apart from each other to act in this way, but they still do it to some extent, more so if the pressure is higher. The CO₂ on Venus, for example, is at such high pressure that its absorption bands are extremely broadened.

Even if the saturation effect were complete, so that the peaks did not get fatter with high concentration, adding more of the greenhouse gas would still increase the temperature of the Earth, for the same reason that adding layers to the layer model also increases the ground temperature. The layers in the layer model are by definition completely saturated in that they absorb all the incoming IR, re-radiating at all frequencies. Adding a second pane of glass to the layer model increases the ground temperature (this is the answer to one of the Exercises). On Venus, even though the absorption of IR by the 70 atm of CO₂ is completely saturated, adding more CO₂ would still increase its temperature further.

The most dilute gases have awesome climate-changing power.

The bottom line is that the more abundant the greenhouse gas is, the less powerful each molecule is in warming the climate. The band saturation for CO₂ makes CO₂ a less potent greenhouse gas, per molecule, than it would be if we had no CO₂ in the air to start with. Revisiting the comparison between CO₂ and methane as greenhouse gases, methane had a disadvantage because its absorption band fell in the suburbs of the Earth-light spectrum, whereas CO₂ fell right downtown. Now, considering the band saturation effect, the advantage shifts to methane because it has much lower concentration. You can see from the jagged edges of the methane peak in Figure 4-3 that the methane absorption band is not saturated. For this reason, in spite of the off-peak location of the methane absorption band, a molecule of methane added to the atmosphere is about 40 times more powerful than is a molecule of CO₂.

For CO₂, the energy intensity I_{out} in units of W/m² goes up proportionally to the log of the CO₂ concentration rather than proportionally to the CO₂ concentration itself (which would be linear in CO₂ concentration). The logarithmic dependence means that you get about the same I_{out} change in W/m² from any doubling of the CO₂ concentration. The radiative effect of going from 10 to 20 ppm pCO₂ is about the same as going from 100 to 200 ppm or 1,000 to 2,000 ppm.

Earth's temperature rises in proportion to the number of CO₂ doublings.

The sensitivities of climate models are often compared as the average equilibrium temperature change from doubling CO₂, a diagnostic that is called the **climate sensitivity** and abbreviated as ΔT_{2x} . It takes several centuries for climate models to respond fully to a change in CO₂, reaching the equilibrium. Most models have ΔT_{2x} between 2° and 5°C, and past climate changes in the real world indicate that the real climate sensitivity of the Earth is probably in this range (Chapter 11).

If you subjected the planet to higher CO₂ concentrations for a few thousand years, there might be changes to the huge ice sheets in Greenland and Antarctica that could lead to even higher global temperatures, so the **Earth-system climate sensitivity** is probably higher than the model ΔT_{2x} values, which do not include such long-term effects. CO₂ emission from fossil fuel combustion will elevate the CO₂ concentration of the atmosphere for hundreds of thousands of years (Chapters 7 and 12), giving plenty of time for the full climate sensitivity of the Earth to express itself.

The climate sensitivity to doubling CO₂ is the first benchmark that climate scientists compare between models.

One can use ΔT_{2x} to calculate the temperature change after a few centuries in response to changing CO₂ using an equation:

$$\Delta T = \Delta T_{2x} \times \frac{\ln \left(\text{new pCO}_2 / \text{original pCO}_2 \right)}{\ln(2)} \quad (4-1)$$




Figure 4-7 The average temperature of the Earth as a function of atmospheric CO₂ concentration and the climate sensitivity parameter, ΔT_{2x} .

where “ln” is the natural logarithm function, the reverse operation of the exponential function e^x . The symbol e denotes a number that has no name other than simply e . We will meet e again in Chapter 5. The exponential function is to raise e to a power of x . If

$$e^x = y$$

then

$$x = \ln(y)$$

Equilibrium temperature changes from changes in CO₂, assuming various ΔT_{2x} values, are shown in Figure 4-7.

The Greenhouse Effect of Gases

The basic idea of the greenhouse effect still works the way it did in the Layer Model (Chapter 3). Figure 4-8 shows the outgoing IR spectrum of an Earth with no CO₂ in the atmosphere. Let us say that this outgoing energy flux is just in balance with incoming sunlight energy.

Now add 1,000 ppm CO₂ to the atmosphere (middle panel), a huge change, so that the effect in the spectra will be easy to see. The energy fluxes in these plots are proportional to the area under the curves. The CO₂ absorption band is like a bite taken out of the curve, decreasing its area, therefore also decreasing the outgoing energy flux. The model reports that I_{out} decreases from 249 W/m² without CO₂ to 223 W/m² with CO₂.

The greenhouse effect drives the surface temperature to balance the energy budget. The online model allows you to change the surface temperature, and by trial and error you can find that warming the Earth up by 8.5°C brings I_{out} back up to about what it was before adding the CO₂. In the bottom panel in Figure 4-8, the effect of the warming can be seen as a shift upward where the arrows are pointing. Especially in the atmospheric window, the intensity of the IR energy flux increases as the Earth warms. Visually, if the spectrum were a pile of ice cream and you wanted to make a volcano, you could scoop out some ice cream from the summit, making the hole where the CO₂ absorption band is, and then spread it back in on the sides of the pile where the atmospheric window is and other parts of the spectrum, so that in the end the overall size of the ice cream pile is the same as it was initially. Adding the CO₂ caused the planet to warm, intensifying the IR energy flux, especially in the atmospheric window, to balance its energy budget.




Figure 4-8 A demonstration of the greenhouse effect of CO₂. In the *top panel*, begin with no CO₂. Assume that the energy budget of the Earth was in balance at a ground temperature of 270 K. In the *middle panel*, add 1,000 ppm CO₂, decreasing the outgoing energy flux. The ground and the atmosphere above it respond by warming up 8.5 K in the *bottom panel*. The total outgoing energy flux is restored to its initial value. The total energy flux is proportional to the area under the curves. CO₂ takes a bite out of the top curve to generate the middle curve, but then the bottom curve bulks up everywhere to compensate.

TAKE-HOME POINTS

- Gases absorb or emit IR light if they vibrate at the frequency of the light and if its vibration has a dipole moment that affects the electric field.
- O_2 and N_2 are not greenhouse gases. All molecules of three or more atoms, including CO_2 , water vapor, and methane, are greenhouse gases.
- A greenhouse gas has a stronger impact on the radiative balance of the Earth if it interacts with light in the middle of the Earthlight spectrum.
- Band saturation: A greenhouse gas at relatively high concentration like CO_2 will be less effective, molecule per molecule, than a dilute gas like methane.

STUDY QUESTIONS

1. Why don't most of the gases in the Earth's atmosphere contribute to the greenhouse effect?
2. Does the band saturation effect prevent global warming from CO_2 ? If not, what effect does it have?
3. Compare and contrast CO_2 and CH_4 as greenhouse gases. Which factor tends to make CO_2 stronger, and which factor tends to make methane stronger? Which gas is stronger overall?

FURTHER READING

Weart, S. *The Discovery of Global Warming*. Cambridge, MA: Harvard University Press, 2003. This is a historical account of the science and the scientists who discovered global warming, including my favorite, Svante Arrhenius, who used the IR spectrum of moonlight, in 1896, to predict that doubling CO_2 would raise global temperature by $3\text{--}6^\circ\text{C}$ (whereas the modern

prediction is $2\text{--}5^\circ\text{C}$). This book includes a good discussion of piecing together the band saturation effect.

IPCC Scientific Assessment 2001, from Cambridge University Press or downloadable from http://www.grida.no/climate/ipcc_tar/. Chapter 6, "Radiative Forcing of Climate Change."

Exercises

Answer these questions using the on-line model at http://understandingtheforecast.org/Projects/infrared_spectrum.html. The model takes CO_2 concentration and other environmental variables as input, and calculates the outgoing IR light spectrum to space, similarly to Figures 4-3, 4-5 and 4-7. The total energy flux from all IR light is listed as part of the model output and was used to construct Figure 4-6.

1. **Methane.** Methane has a current concentration of 1.7 ppm in the atmosphere and its doubling at a faster rate than is CO_2 .
 - a) Would an additional 10 ppm of methane in the atmosphere have a larger or smaller impact on the outgoing IR flux than an additional 10 ppm of CO_2 at current concentrations?
 - b) Where in the spectrum does methane absorb? What concentration does it take to begin to saturate the absorption in this band? Explain what you are looking at to judge when the gas is saturated.

- c) Would a doubling of methane have as great an impact on the heat balance as a doubling of CO_2 ?
- d) What is the “equivalent CO_2 ” of doubling atmospheric methane? That is to say, how many ppm of CO_2 would lead to the same change in outgoing IR radiation energy flux as doubling methane? What is the ratio of ppm CO_2 change to ppm methane change?

2. CO_2

- a) Is the direct effect of increasing CO_2 on the energy output at the top of the atmosphere larger in high latitudes or in the tropics?
- b) Set pCO_2 to an absurdly high value of 10,000 ppm. You will see a spike in the CO_2 absorption band. What temperature is this light coming from? Where in the atmosphere do you think this comes from?

Now turn on clouds and run the model again. Explain what you see. Why, at night, is it warmer when there are clouds?

3. Water vapor. Our theory of climate presumes that an increase in the temperature at ground level will lead to an increase in the outgoing IR energy flux at the top of the atmosphere.

- a) How much extra outgoing IR would you get by raising the temperature of the ground by 5°C ? What effect does the ground temperature have on the shape of the outgoing IR spectrum and why?
- b) More water can evaporate into warm air than into cool air. By setting the model to hold the water vapor at constant relative humidity rather than constant vapor pressure (the default), calculate again the change in outgoing IR energy flux that accompanies a 5°C temperature increase. Is it higher or lower? Does water vapor make the Earth more sensitive to CO_2 increases or less sensitive?
- c) Now see this effect in another way. Starting from a base case, record the total outgoing IR flux. Now increase pCO_2 to 550 ppm. The temperature in the model stays the same (it's just how the model is written), so the IR flux goes down. Now, using the constant vapor pressure of water option, increase the temperature offset until you get the original IR flux back again. What is the change in T required? Now repeat the calculation but at constant relative humidity. Does the increase in CO_2 drive a bigger or smaller temperature change?

How the greenhouse effect works within the temperature structure of Earth's atmosphere

The Greenhouse Effect Is Powered by the Lapse Rate

For the greenhouse effect to work, it has to be colder up high in the atmosphere than it is on the ground. It is colder up high in the Layer Model, and in the real atmosphere, too, so the greenhouse effect works in both settings. However, the temperature structure of the real atmosphere is controlled by processes that are not in the Layer Model, in particular a process called moist convection, which supersedes the radiation energy fluxes in the Layer Model. Moist convection is controlled by the interaction of several pieces, which may seem to be unrelated, until they are brought together at the end of the chapter.

In the Layer Model, the pane of glass is colder than the ground, and it absorbs the intense infrared (IR) energy flow from the ground and reradiates weaker IR to space. This perturbation to the energy flow results in greenhouse effect warming.

Now imagine installing a heat-exchanger system running between the glass and the ground, that carries enough heat that it equalizes the temperatures of the glass and the ground when you dial it up all the way (Figure 5-1). Now the ground has the same temperature as the glass, which as always is at skin temperature. The ground is no warmer than it would have been if there had been no glass; the greenhouse effect has vanished. The intensity of the IR energy that the glass emits is the same as the IR energy it absorbs from the ground, because the temperatures of the ground and the glass are the same. The glass does not change the amount of outgoing IR, so it does not make the ground warmer. The diligent student may have also discovered this phenomenon in the exercise from Chapter 3 on nuclear winter.

No temperature contrast, no greenhouse effect.

Now dial the heat exchanger down a bit, allowing a temperature difference to build up between the glass and the ground. After enough time has passed for the energy flows and temperatures to reach their steady state, the temperature of the glass will return to the value it was before, because the glass is the outermost or “skin” layer (Chapter 3). The energy flux from the skin layer has to balance incoming sunlight, fixing its temperature at the “skin” temperature. If the glass temperature is the same as before, the ground temperature must be warmer than it was before. It turns out you can dial the greenhouse effect up and down by dialing up and down the temperature difference between the skin and the ground.

Atmospheric scientists call the change in temperature of the air with altitude the **lapse rate**; in the lower part of Earth's atmosphere where the greenhouse effect operates, it is about 6°C colder for each kilometer of altitude. The lower part of the atmosphere is called the **troposphere**




Figure 5-1 The Layer Model from Chapter 3 with an added heat exchanger, capable of partially or completely equalizing the temperatures of the two layers.

(Figure 5-2), and it contains 90% of the air and all of the weather. The air temperature reaches its coldest point at the **tropopause**, a boundary of air about 10 km high on average, about where commercial airplanes fly. Above this, it gets warmer with altitude in the **stratosphere** because of ozone which absorbs ultraviolet radiation from the sun (Chapter 10), and higher still are the mesosphere and the exosphere, which do not affect the climate story very much.




Figure 5-2 Typical temperatures and pressures of the atmosphere as a function of altitude in the tropics.




Figure 5-3 A demonstration of the effect of the lapse rate on the strength of the greenhouse effect. If we increase the greenhouse gas concentration of the atmosphere, the infrared radiation to space will originate from a higher altitude (labeled *Skin altitude*), but the skin temperature at the skin altitude will eventually have to be the same as the temperature at the old skin altitude. The increase in skin altitude increases the ground temperature. If the temperature of the atmosphere were the same at all altitudes, then raising the skin altitude would have no impact on ground temperature.

The idea of a skin layer in the atmosphere is fuzzier than it is in the layer model, but it is still useful.

In a column of air, instead of two layers, there is air at a range of temperatures (Figure 5-3). IR radiation comes up from the ground, and some of it might make it all the way to space unhindered; the rest is absorbed and re-emitted by greenhouse gases and clouds (Chapter 4). There is no single place where all the IR radiation escapes to space from, like the skin layer in the Layer Model, but it is useful to define a **skin altitude** for the air column as some kind of average altitude from which the IR escapes to space.

In solving the layer model for the temperatures of all the atmospheric layers and the ground, it was convenient to start from the outer skin layer and work down toward the ground. Let us take that approach again in thinking about the column of air. If the skin altitude were 5 km high and the lapse rate were set by some physics of the atmosphere to be 6°C per kilometer, then the temperature of the ground would be

$$T_{\text{ground}} = T_{\text{skin}} + \frac{6^{\circ}\text{C}}{\text{km}} \cdot 5 \text{ km}$$

On the temperature profile plot (Figure 5-3), this looks like a line coming downward from the skin altitude and temperature, following the slope of the lapse rate, until it intersects the ground.

If you put more greenhouse gases in the air, more of the IR will come from colder high-altitude air, and the skin altitude, which is the average over all frequencies of IR, moves a bit higher up. Right after you add the greenhouse gas, before the planet has had time to warm up, the temperature at the new skin altitude will be colder than it was at the old skin altitude because

the new skin altitude is higher than the old one. But as in the Layer Model, after some time the temperature at the skin altitude has to be just warm enough to balance the energy influx from sunlight. The planet warms, and the entire column of air warms with it. Eventually, it warms until the temperature at the new skin altitude is the same as the temperature was at the old skin altitude.

More CO₂ → Higher skin altitude → Warmer ground

Beginning from this new “anchor point,” the temperature of the ground can be determined by drawing a line from the skin altitude toward the ground, again following the slope of the lapse rate. The change in temperature at the ground ΔT_{ground} that you would get from raising the skin altitude, Δz_{skin} , can be calculated as

$$\Delta T_{\text{ground}} = \Delta z_{\text{skin}}[\text{km}] \cdot \frac{6^\circ\text{C}}{\text{km}}$$

Now imagine how it would work if the lapse rate were different from 6°C per kilometer. A change in the skin altitude would have a different effect on the temperature of the ground. For example, if the physics of the atmosphere were to decide somehow that the temperature should be uniform with altitude, like when we dialed up the heat exchanger in the Layer Model above, then it would make no difference how high the skin altitude was:

$$\Delta T = \Delta z_{\text{skin}}[\text{km}] \cdot \frac{0^\circ\text{C}}{\text{km}} = 0$$

The temperature of the ground would be the same as the skin temperature no matter what. Just like for the Layer Model, you can dial up and down the strength of the greenhouse effect by changing the lapse rate in the atmosphere.

Pressure as a Function of Altitude

The pressure in the atmosphere depends primarily on the weight of the air over your head. As you climb a mountain, the pressure decreases because you are leaving atmosphere behind; the air beneath you does not contribute to the pressure you feel.

In the ocean, every 10 m deeper you go will increase the pressure by about 1 atmosphere (atm). Each 10-m depth interval increases the pressure by about the same 1 atm: descending from 30 to 40 m would increase the pressure by the same 1 atm as going from 0 to 10 m. The pressure can be calculated as

$$P = 1 \text{ atm} + \frac{-1 \text{ atm}}{10 \text{ m}} \cdot z[\text{m}]$$

where the variable z is the vertical position, with positive numbers upward, as before, so a negative height is a depth.

The density of water does not change much with pressure, so the weight of the water overhead (the pressure) depends linearly on the water depth.

Because pressure is just a constant number multiplied by the depth, a graph of pressure versus depth would be a straight line (Figure 5-4), and the pressure is said to be **linear** with depth.

The pressure in the atmosphere is **nonlinear** with altitude in that a climb of 1 m at sea level changes the pressure much more than 1 m up at the tropopause. The equation to describe




Figure 5-4 Pressure increases linearly with depth in water because water is incompressible.

pressure as a function of height in the atmosphere is based on the **exponential function**, which is a number called e raised to a power. The value of e is approximately 2.71828.

The density of a gas increases with pressure, so the weight of air overhead (the pressure) follows an exponential function.

The exponential function comes up time and time again in the natural sciences. Population growth and radioactive decay are two examples (Figure 5-5). In each case, the rate of change of the variable is proportional to the value of the variable itself. Population growth is driven by the number of babies born, which depends on the number of potential parents to beget them:

$$\text{Population}(t) = \text{Population}(0) \cdot e^{kt}$$

In this expression the exponent is a population growth rate in units of fraction per time, like 3% per year might be a typical population growth rate. This is an example of a positive exponential function; it is essentially an explosion, either a slow one or a fast one.

Radioactive material decays at some well-determined rate. Each decay releases some form of radiation, such as can be counted with Geiger counter, but the radioactivity decreases with time as the substance decays away. Here the exponential function is negative, decaying rather than growing:

$$\text{Counts}(t) = \text{Counts}(0) \cdot e^{-kt}$$

The exponential function comes up when a quantity grows or shrinks in proportion to its current value.




Figure 5-5 The exponential functions e^x and e^{-x} . For e^x , the growth rate of the function is proportional to its value. Examples of this type of behavior include interest on a bank account and population growth. For e^{-x} , the value of the function decays proportionally to its value. Radioactive decay does this, as does pressure as a function of altitude in the atmosphere.

A decaying exponential approaches zero **asymptotically**, meaning that it gets closer and closer but never actually gets there. In the radioactive decay equation, the scaling factor in the exponential could be converted into an **e -folding time scale**, as

$$e\text{-folding time scale} = 1/k$$

This quantity is similar to a half-life but rather than the amount of time to decay to half the initial quantity, the e -folding time scale tells how long until 37%, $1/e$, of the original is left. It takes 44% longer to decay to a fraction of $1/e$ than it does to $1/2$ of the original number of atoms.

The atmospheric pressure decreases as a function of altitude according to an exponential decay type of equation

$$P(z) = 1 \text{ atm} \cdot e^{\frac{-z[\text{km}]}{8 \text{ km}}}$$

shown in Figure 5-1. The height z is zero at sea level, leaving us with e^0 , which equals 1, so the pressure at sea level is 1 atm. At an altitude of 8 km, pressure is lower than at sea level by a factor of e^{-1} or $1/e$, about 37%. We call that altitude the **e -folding height**.

Most of the mass of the atmosphere is contained in the e -folding height. In fact, if the entire atmosphere were at 1 atm of pressure, instead of smeared out in a decaying exponential function, it would fit into one e -folding height exactly. So in one sense you could say that the atmosphere is 8 km thick, even though, mathematically, according to the exponential function, the pressure in the atmosphere never really gets to zero, no matter how far away you get.

From the appearance of the exponential function in the equation for pressure, you might guess that the rate of change of pressure with altitude must depend on the pressure itself in some way. This would be astute. The rate of change of pressure depends on pressure because at high pressure, gas is compressed, and so a climb of 1 m through gas at high pressure would rise above more molecules of gas than would a climb through a meter of gas at low pressure.




Figure 5-6 A wall made of compressible bricks, to demonstrate why pressure changes rapidly at the bottom of the atmosphere and slowly at the top because air in the atmosphere is compressible (*left*). For water, which is incompressible, pressure increases linearly with depth.

Imagine a wall made of compressible bricks (Figure 5-6). A row of bricks is thinner at the bottom of the wall because they are compressed. Batman climbing up the wall would pass more bricks per step at the bottom than at the top. For incompressible bricks (the normal kind), the rows are all the same height and the mass of wall above you is a linear function of height.

The bottom line is that gas expands when depressurized, and because of this, atmospheric pressure decreases nonlinearly with altitude, following an exponential decay function.

Expansion, Compression, and Heat

Here begins the second apparently unrelated thread of the convection story. If we compress a gas, its temperature goes up. This occurs even if we do not allow any heat to enter the gas or leave it, say if we had gas inside an insulated piston that we compress or expand. The condition that we are describing, a closed system with no heat coming in or out, is called **adiabatic**. If gas is compressed adiabatically, it warms up. If you ever let the air out of a bicycle tire by holding the little valve open with your thumb, you may have noticed a chill on your thumbnail. The gas expands as it flows from the high-pressure tire into the lower-pressure atmosphere, and because it expands, it cools.

If you squeeze air, does it not heat up?




Figure 5-7 The temperature decrease with altitude in the atmosphere is called the lapse rate. If the atmosphere contained no water vapor, the lapse rate would follow the dry adiabat. As is, most of the atmosphere is closer to the moist adiabat, especially in the tropics, where convection is most important.

When you compress a gas, its temperature rises. It takes work to compress a gas. You have to squeeze the piston to compress it. Your muscles push against molecules bouncing against the head of the piston. That work energy is transformed into heat, the bouncing-around energy of the molecules of gas. The reverse situation, cooling while expanding, is a little less intuitive to envision, but when the piston moves to expand the gas, it is taking work energy from the gas. You could use the power of the piston to do some work, like turn a driveshaft. The energy to do that work comes from the thermal energy of the molecules of the gas, so the gas cools as it expands.

Figure 5-7 shows the temperature change that a parcel of dry surface air would experience if it were carried aloft. We'll worry about water vapor next, but for now we are looking at the line labeled **Dry adiabat**.

Water Vapor and Latent Heat

Here comes seemingly unrelated thread number three. Water molecules can exist together in any of three **phases**: gas, liquid, or solid. A transformation from liquid or solid to gas requires an input of a considerable amount of energy. One could write a chemical reaction for water as



If you have ever burned your skin with steam from a teapot, this reaction will have meaning for you. Steam from a teapot is probably at the boiling temperature of water, 100°C or 212°F . This is a temperature you might set for your oven to warm bread; you can easily blow the air from the oven on your hand when you reach in for the bread without burning yourself. Steam burns, not because its temperature is so high (what is called its **sensible heat** because a thermometer can sense it), but because the vapor deposits its heat when it condenses on your skin. The heat

that the steam is carrying is called **latent heat**. You charge up an air parcel with latent heat when you evaporate water into it, and you get the heat back again when the water condenses. A thermometer does not measure the latent heat content of an air parcel unless the heat is released by condensation.

If you set a tray of water or ice in a closed box full of dry air, some of the water molecules are going to evaporate, becoming a gas. Actually, the air in the box is almost irrelevant to the water. Almost the same number of water molecules would evaporate into a box full of vacuum (no air). A chemical way to say that is that, to a first approximation, the partial pressure of water vapor does not depend on how much oxygen or nitrogen is there; it only depends on the temperature inside the box.

When water is in equilibrium between liquid and vapor, it's called saturated, or 100% relative humidity.

When the temperature is high, more molecules will have the energy required to jump into the vapor phase, and the **equilibrium vapor pressure** of water will be high. In cold conditions, the equilibrium vapor pressure will be lower (Figure 5-8). If the amount of water vapor actually present is lower than the equilibrium value, it is what is called **undersaturated**, and liquid water, if available, will tend to evaporate. If water vapor is **supersaturated**, meaning the vapor pressure is higher than equilibrium, then water will tend to condense, perhaps into raindrops or snowflakes, or it could wet the walls of the box. For water vapor, the equilibrium pressure is also referred to as 100% **relative humidity**. More on water vapor to come in Chapter 7.




Figure 5-8 Water has a greater tendency to evaporate at higher temperature than at cooler temperatures. The graph shows what a chemist would call the equilibrium vapor pressure of water and a meteorologist would call 100% relative humidity. Air in the real world can have less or sometimes slightly more water vapor than this, but if the water vapor concentration gets too much higher than this curve, the water will condense into water droplets or snow.

Convection

Now we are equipped to think about convection in the atmosphere. **Convection** occurs when you heat a fluid or gas from below or cool it from above. If you heat it from below, the fluid at the bottom gets warmer than the fluid overlying it. Warmer molecules bounce around more energetically, pushing outward more strongly; as a result, the fluid expands somewhat as its temperature increases. As it expands, its density (mass per volume) decreases. The result of heating from below is warm buoyant fluid underneath colder denser fluid, an unstable situation that causes the fluid column to turn over, the warm fluid rising to the top.

Consider a pan of water on a stove. Figure 5-9 shows the temperature of the water as a function of the height in the pan, a temperature profile. First mix the water in the pan. When the water is well mixed, it will all have the same temperature. That sounds obvious, but we shall find in a moment that a well-mixed column of air does not all have the same temperature. A well-mixed, uniform-temperature water column is called **statically stable**; it is not stratified, but it is not unstable either.

Next turn on the burner, warming and expanding the water at the bottom of the pan. Buoyant water at the bottom tends to rise; this situation is called **convectively unstable**. The buoyant water from the bottom could potentially rise to the top like the lava in a lava lamp, in which case the water column would be **stratified**.




Figure 5-9 The effect of convection on the temperature in a pan of water on a stove. At *top left*, the water in the pan is well mixed, the same temperature throughout. After the burner is turned on, the water at the bottom of the pan warms, which makes it less dense than the water above it (convectively unstable, *top right*). Either this water can rise to the top without mixing (*bottom left*) or it can mix with the water above it, raising the temperature of the entire pan of water (*bottom right*). The atmosphere tends to mix when it convects, as in the *lower right* example.




Figure 5-10 Convection in a compressible fluid such as air. On the left is the statically stable configuration, analogous to the *upper left panel* in Figure 5-9. The temperature decreases with altitude because the gas expands and therefore cools. This temperature profile is what you would see in a well-mixed column of air. On the right is the effect of heating from below, analogous to the *top right panel* in Figure 5-9.

The temperature profile in a convecting pan of water evolves through a series of well-mixed (statically stable) states.

Alternatively, the rising warm water could mix with the rest of the water, as generally occurs to a large extent in the atmosphere and ocean. In this case we end up after the convection with a second statically stable profile, at a higher temperature than the initial one.

Convection in a column of air is like that in a column of water, but with a few new twists added. As with the water, start by mixing everything together (Figure 5-10). After mixing, no air parcel will be different than any other because it is all the same air. The column of air will be statically stable because that is what mixing does.

Unlike the incompressible water, however, the air is not all the same temperature. The pressure is higher at the bottom of the column because of the weight of the air, so the air at the bottom has been compressed and it has heated up. If a warm air parcel from the bottom were raised up to the top, it would expand and cool, but because the air is well mixed, the moving air will always find itself at the same temperature as the rest of the air in the column. The parcel was neutrally buoyant near the bottom of the column, and it is neutrally buoyant all the way up. This is what static stability looks like in a column of compressible air.

A well-mixed (statically stable) column of gas (a compressible fluid) will not have the same temperature throughout. It will be warmer on the bottom.

Ocean water is slightly compressible, so its temperature changes somewhat with depth in a well-mixed column, just like air does, but the effect is much smaller for water, just a few degrees Centigrade change between the surface and the bottom of the ocean.

Convection in the atmosphere is driven by sunlight hitting the ground, warming the air at the bottom of the air column. The warmed air parcel from the ground finds itself warmer than the air around it, and so it begins to rise. As it does, it expands and cools, but throughout the ascent the air parcel remains warmer and lighter than the air surrounding it. If it does not mix with surrounding air on the way up, it can rise all the way to the top of the column.

If the rising plume of air mixes with the surrounding air as it goes, as it usually does in the atmosphere, the whole air column would warm up uniformly, just as it did in the column of

water. And, just like the water, the temperature profile in the column would always be statically stable as it rose. For the atmosphere, the temperature profile follows the curve labeled Dry adiabat in Figure 5-7.

Moist Convection

Convection in the real atmosphere has one more ingredient, and that is latent heat from water vapor. Imagine a source of water vapor at the Earth's surface so that the air above the ground has a relative humidity of 100%. If this parcel rises up in the column due to convection, it will decompress and cool.

As the temperature drops, the equilibrium amount of water vapor decreases (Figure 5-8). The decreasing saturation drives an increase in the relative humidity value (since the saturation value is what it's relative to), which exceeds 100% as the air cools. The supersaturation drives water to condense into droplets or ice. The story of cloud formation in climate is not a simple one and does not end here but continues in Chapter 7.

As the water condenses, it releases its latent heat, just like the steam from the teapot burning your skin. It takes heat to make vapor from liquid or gas, and when you re-form the liquid or gas, you get the heat back. The latent heat warms the air parcel up a bit, counteracting some of its decompression cooling (Figure 5-11). The water vapor changes the temperature of the air; in fact, it systematically changes the lapse rate (temperature decrease with altitude). To distinguish the two, atmospheric scientists talk about **dry convection** and **moist convection**, and **dry** and **moist adiabatic lapse rates**.

The latent heat in water vapor drives most of the drama in our weather.

If there were no convection, only radiation like in the Layer Model, the lapse rate of the atmosphere would be about 16°C per kilometer of altitude. The atmosphere does not get there because that state would be unstable to convection, like a pan of water with warmer water underneath cooler water. There are some stars that have radiation energy fluxes like in the Layer Model but have no convection. They are said to be in a state of **radiative equilibrium**. But the atmosphere of the Earth convects, finding a state called **radiative convective equilibrium**.

An atmosphere undergoing dry convection has a lapse rate of about 10°C temperature change per kilometer of altitude. When you add the latent heat in moist convection, the lapse




Figure 5-11 The effect of water on the temperature profile in the atmosphere. Air cools as it rises, as in Figure 5-10. The cooling drives water vapor to condense into droplets, releasing its latent heat to the air. The moist adiabat therefore cools less with altitude than does the dry adiabat.

rate decreases to about 6°C per kilometer. This is very close to the lapse rate of the real atmosphere. As the next chapter continues to show, the circulation of the atmosphere is rather complicated, but the process of moist convection is a good first-order description of why the temperatures in the atmosphere are stacked the way they are.

The question is important because it is possible that the lapse rate of the atmosphere could be different in a changing climate. This would be one of many feedbacks in the climate system (Chapter 7).

TAKE-HOME POINTS

- Air in the upper troposphere is colder than air at the ground because of the process of moist convection. The process includes the following:
 - a) Convection is driven by sunlight heating the air near the ground.
 - b) The air rises and cools because it expands.
 - c) Water vapor condenses, releasing heat as the air rises.
- The moist convecting air gets colder with altitude, but not as much as if it were dry.

If the air did not get colder with altitude at all, there would be no greenhouse effect.

STUDY QUESTIONS

1. Explain why a well-mixed column of gas is warmer on the bottom.
2. If the atmosphere got warmer with altitude, instead of cooler as it does in Earth's troposphere, what would be the climate impact of greenhouse gases?
3. How can the skin altitude of the atmosphere be defined?
4. Why does the profile of atmospheric pressure with altitude in the atmosphere follow an exponential function?
5. A parcel of air rises convectively in the atmosphere, losing its water as rain and getting cooler following the moist lapse rate. It gets sucked into an airplane, where it is recompressed to the pressure at the ground without replacing the water. Is this air warmer or colder than air at the ground? What is its relative humidity like?

FURTHER READING

Frederick, J. E. *Principles of Atmospheric Science*. Sudbury, Mass.: Jones and Bartlett Publishing, 2007.

Exercises

1. **Lapse Rate.** Use the online full-spectrum radiation model at http://understandingtheforecast.org/full_spectrum.html. Adjust the lapse rate in the model and document its impact on the equilibrium temperature of the ground.
2. **Skin Altitude.** Answer this question using the online IR radiation model.
 - a) Run the model in some configuration without clouds and with present-day pCO_2 . Compute σT^4 using the ground temperature to estimate the heat flux that you would get if there were no atmosphere. The value of σ is $5.67 \cdot 10^8 \text{ W}/(\text{m}^2 \text{ K}^4)$. Is the model heat flux at the top of the atmosphere higher or lower than the heat flux you calculated at the ground?

- b) Now calculate the “apparent” temperature at the top of the atmosphere by taking the heat flux from the model and computing a temperature from it using σT^4 . What is that temperature, and how does it compare with the temperatures at the ground and at the tropopause? Assuming a lapse rate of 6 K/km and using the ground temperature from the model, what altitude would this be?
- c) Double CO_2 and repeat the calculation. How much higher is the skin altitude with doubled CO_2 ?
- d) Put CO_2 back at today’s value, and add cirrus clouds. Repeat the calculation again. Does the cloud or the CO_2 have the greatest effect on the “skin altitude”?

How the weather affects the climate

Chaos

Weather forecasts have really improved in the last few decades, but it is still fairly dicey to say what the weather will be 10 days from now. And yet climate science claims to be able to forecast the climate 100 years from now. Perhaps an explanation is in order.

It is indeed tricky to forecast the weather too far in advance because weather is **chaotic**. To a scientist, the word chaotic means an extreme sensitivity to initial conditions, so that small differences between two states tend to amplify, and the states diverge from each other. This behavior is called the **butterfly effect**, the reference being to a puff of air from a butterfly's wing eventually resulting in a giant storm somewhere that would not have happened if the butterfly had never existed.

The butterfly effect was first observed in a weather simulation model. The model computer run stopped, and the researcher Edward Lorenz restarted it by typing in the values of model variables like temperatures and wind speeds, but with small seemingly insignificant changes, like rounding errors. It did not take long for the restarted simulation to diverge completely from the results of the initial simulation, because of the chaotic sensitivity of Lorenz's model to its exact starting point.

Weather is chaotic, but climate generally is not.

Weather forecasts rely on computer models that are very similar to climate models. The forecast model runs are started from meteorological data and past model results. Small imperfections in the initial conditions, and in the model, cause the model weather to diverge from the real weather, so that by about 10 days the model prediction becomes meaningless.

One way to cope with this situation is to run the model lots of times with tiny variations in initial conditions—an **ensemble** of model runs. Then you can see something of the range of what is possible. It does not fix the problem exactly, but it does result in a more reliable weather forecast.

The **climate** is defined as some time-average of the weather. One could speak of a **climatological** January, which would be the average of many Januaries, say 10 or 20 Januaries. Forecasting climate is not as difficult as forecasting weather because it does not matter if the rain you predict comes on Tuesday or Thursday. The butterfly cannot foil our climate prediction with her storm if the model gets the large-scale tendency for storminess correct. If the model does not follow a storm trajectory this time, it will next time; the butterflies average out. The weather is chaotic, but the climate is not.

The butterfly is able to change the specifics of a particular storm, but it is not able to change the general storminess of the climate. The statistics of the weather events in the climate

model ought to replicate those of the real climate, even if the exact storms are not the same between the model and reality.

Averaging

The layer model represents Earth's climate on a grand average: the warm places and the cold places, the summers and the winters, day and night. The Layer Model also assumes that the radiative energy budget is completely balanced. This is a reasonable assumption for the long-term global average because there is nowhere else for the heat energy to go; what comes in must go out.

However, in the real world, some places are hotter than others. Very warm and very cold places on Earth are significantly different from average places in terms of how their energy budgets work. Also, the real radiative energy budget for some location and time is probably wildly out of balance. The sun shines only in the daytime, for one thing; energy is dumped onto the Earth in the daytime and radiates away through the night. There is also an imbalance through the seasonal cycle because it takes a buildup of heat every spring to warm things up: the air, land surface, and especially the water.

Is an eternal, unchanging, averaging model a reasonable one for a world that is usually locally out of balance? Can you construct the average of the whole system by using averages of the pieces of the system? Or will averaging change the answer?

There could be a problem with averaging infrared (IR) energy fluxes, for example, because they are a **nonlinear** function of temperature, which is to say that a plot of one variable versus the other does not trace a straight line (Figure 6-1). Algebraically, one can look at the equation and see that the light flux is proportional to temperature to the fourth power, and not to the first power, which would be linear.

$$I \left[\frac{\text{W}}{\text{m}^2} \right] = \varepsilon [\text{unitless}] \sigma \left[\frac{\text{W}}{\text{m}^2 \text{ K}^4} \right] T [\text{K}]^4$$

Let's estimate the average energy flux of a planet that had two sides, one at 300 K (rather like Earth) and the other at 100 K (much colder than anywhere on Earth). The outgoing energy flux from the cold side would be about 6 W/m², and from the warm side, 459 W/m². The average energy flux would be 232 W/m².

It may be OK to take the average of a linear system, but averaging a nonlinear system may introduce a bias.

Now average the temperatures; run the average T through $\varepsilon \sigma T^4$. The average temperature is 200 K, and the calculated energy flux is 91 W/m². The averaged answer is off by more than a factor of 2. You can see the effect graphically as the difference between the straight line and the curve in Figure 6-1.

The terrestrially normal range of temperatures spans a narrow enough range that it looks much straighter than in Figure 6-1, which spans an absurdly wide temperature range. An averaging model would work better for realistic temperature extremes. Still, the variations in temperature could have an important impact on the behavior of the system, and we certainly would not want to neglect it in the global warming forecast.

There are many other nonlinearities in the climate system, pairs of variables that are related to each other in nonlinear ways. For example, many effects of the wind, such as exchanging heat with the ocean, depend nonlinearly on wind speed.




Figure 6-1 An example of how averaging can bias the results of a nonlinear system such as blackbody radiation energy flux $\epsilon\sigma T^4$. In the *top panel*, we average across a huge temperature range, and the flux we get from combining areas of 100 and 300 K is very different from what we would expect if we averaged the temperature first to 200 K and then computed the answer. In the *bottom panel*, we see that over the temperature range of normal Earth conditions, the blackbody radiation energy flux is closer to linear, so averaging would be less of a problem than in the *top panel*.

The Fluctuating Heat Budget

The heat input to the real Earth is not stable and unchanging like the Layer Model; instead, it fluctuates rather widely. Solar energy comes in only during the day, for example, whereas IR energy loss to space continues around the clock (Figure 6-2). The heat input also varies over the seasonal cycle and from place to place on the Earth's surface.

Earth's heat budget is only in balance on average, but at any given instant of spot on the Earth it is usually out of balance.




Figure 6-2 The surface of the Earth receives only incoming solar radiation during the daytime (*heavy solid line*), but it radiates infrared light all the time (*thin solid line*). The energy budget for this location (*dashed line*) is in balance only when it is averaged over 24 hours.

The seasons are caused by the tilt of the Earth relative to its orbit around the sun (Figure 6-3). In the winter hemisphere, the days are shorter and the sun is at a lower, more oblique angle in the sky. Added up over a day, the winter hemisphere has less sunlight energy. The seasons are opposite each other in the two hemispheres because if one hemisphere is tilted toward the sun, the other will be tilted away.

Note that the seasons are not caused by the relative distances of any place on Earth from the sun. The North Pole is not really that much farther from the sun than the equator is, because the size of the Earth is much smaller than the distance of the Earth from the sun.

Neither are the seasons caused by changes in the distance of the whole Earth to the sun. It is true that Earth's orbit is elliptical, rather than circular, so the distance to the sun does vary throughout the year. But in its current orbital configuration, the Earth is closer to the sun in January than in July, which is the wrong way to explain winters in the northern hemisphere. If the distance from the sun caused the seasons, the whole world would get cold and warm at the same time.

The effect of the tilt and eccentricity of the Earth's orbit on sunlight intensity is shown in Figure 6-4. The axes in this figure are day of the year in the x -direction and latitude in the y -direction. The contours show the intensity of sunlight at the top of the atmosphere averaged over 24 hours. Any location at some latitude, say 42 degrees north, which goes through Chicago,




Figure 6-3 The Earth's tilt is responsible for the seasons. This is southern hemisphere summer, northern hemisphere winter.




Figure 6-4 The Earth's tilt determines how much heat the surface receives from the sun each day as a function of latitude (*vertical axis*) and time of year (*horizontal axis*).

Barcelona, and Vladivostok, will get the same solar influx at the top of the atmosphere. A map of the solar intensity would look like a bunch of horizontal stripes. Figure 6-4 shows the two dimensions over which the solar influx does vary: latitude and day of the year.

Earth's seasons are caused by tilt of the poles relative to the orbit,
not by its distance from the sun.

The beginning of the year in this plot is southern hemisphere summer, and there are large numbers, representing high energy flux in the south, and small numbers in the north. Northern hemisphere summer comes in the middle of the plot, centered on about day 180. Now there is high energy flux in the north and lower in the south. The sunlight intensity pattern arises by two mechanisms. First, the intensity of sunlight per square meter (m²) of ground area is greater at noon in the summer than in the winter because the ground is at less of an angle to the sun in the summer (Figure 6-3). Second, days are longer in summer, and this increases the 24-hour average energy influx.

It is interesting to note that the highest daily-average energy fluxes on Earth are at the North or South Pole during northern or southern summer. The sun never sets in midsummer at the pole; it just whirls around in a circle over the horizon. In winter the pole gets no sunlight for months on end. The poles do not turn into tropical garden spots in summer because to some extent the temperatures average out the seasonal changes in heat forcing.

Thermal inertia tends to damp out the temperature swings of the day-night cycle, the seasonal cycle, and any global warming temperature trends as well. The liquid ocean in particular has a tremendous capacity to absorb and release heat to the atmosphere. In soils and on the land surface, the seasonal cycle of temperature reaches only a meter or two down because the diffusion of heat through solids is relatively slow. When you cool the water surface, turbulent mixing and convection quickly spread that cooler water into the water column, typically 100 m or more. For this reason, the seasonal cycle is much more intense in the middle of large continents than it is in maritime areas impacted by the temperature stabilizing effects of liquid water.




Figure 6-5 The energy budget between incoming solar and outgoing infrared radiation does not balance locally because heat is transported on Earth by winds and currents. The equator receives more solar energy than it radiates as infrared, whereas high latitudes lose more infrared than they receive from sunlight.

Even when averaging out the seasonal cycle, the radiative energy budget of a single spot on the surface of the Earth is still probably way out of balance, because heat energy is redistributed around the Earth's surface by wind and water currents (Figure 6-5). There is a net influx of heat in the tropics as sunlight brings in energy faster than the outgoing IR. It does not get hot enough in the tropics for outgoing IR to balance the incoming sunlight locally. Heat is carried into the cooler high latitudes by warm winds and ocean currents. In high latitudes, the Earth vents the excess tropical heat as excess radiative heat loss to space.

The upshot of this chapter is that a climate model to forecast global warming will have to simulate the weather: all of the fluctuations and imbalances in the energy budget and the way the Earth's climate responds, by storing heat or moving it around. The Layer Model will not do.

The Coriolis Acceleration

Flows in the atmosphere or ocean differ from those in a sink or bathtub because they persist long enough that they are steered by the rotation of the Earth. When water runs down the drain of a bathtub or sink, it takes only a few seconds for the swirling pattern at the drain to establish itself, and the Earth has only rotated a negligible amount in that time. So it is not true that water has much tendency to spiral down the drain in opposite directions in the two hemispheres! We cannot feel the Earth's rotation in our everyday motions either. But if the flow down the drain persisted for several days, the rotation of the Earth would start to become significant to steering the flow.

Imagine two kids sitting opposite each other on a merry-go-round at a playground (Figure 6-6), tossing a ball back and forth while spinning. The actual rotation of the Earth is negligible here because the ball takes less than a few days to get where it is going, but the merry-go-round itself is rotating quickly enough to move as the ball flies from one side to another. The thrower (T) tosses the ball toward the catcher (C), but before it gets there, the merry-go-round has rotated, and the catcher misses the ball. Viewing the action from the ground (the nonspinning frame of reference), the ball flies straight, following Newton's laws of motion. The rotation of the merry-go-round underneath the ball does not actually affect the trajectory of the ball.

A person on Earth looking at a map of the weather or the Gulf Stream is like the kids on the merry-go-round trying to make sense of the apparent trajectory of the ball. It is inconvenient to




Figure 6-6 An illustration of the Coriolis acceleration on a playground merry-go-round. The thrower (T) tosses the ball toward the catcher (C). Before the ball gets to the catcher, the merry-go-round rotates. If we observe this from the point of view of the children on the merry-go-round, it appears as if the ball veers to the right as it moves. The Coriolis acceleration is the kids' attempt to explain the motion of the ball, as though the merry-go-round were not spinning.

try to visualize things from the nonrotating frame. What they see is that, on the merry-go-round, the ball curved to the right.

Newton's laws of motion can be “hacked” to get them to work in the rotating frame by adding a fake force called the **Coriolis acceleration**. This fudge factor has the same effect as the rotation of the Earth that is being ignored. On the merry-go-round, the fake force should change the direction of the ball without changing its speed. To do this, the fake force must be applied at exactly 90 degrees perpendicular to the direction that the object is moving. It also turns out that the force must be stronger the faster the object is moving, because it would take a stronger force to change the trajectory of a faster object. In Figure 6-6, the Coriolis fake force would have to be directed to the right of the direction of motion so that it would turn the trajectory of the ball to the right. On Earth, the Coriolis force goes to the right of the direction of motion in the northern hemisphere and to the left in the southern hemisphere.

The Coriolis force is a fudge factor, allowing us to ignore, in the equations of motion, the fact that the Earth is rotating.

Now leap from a flat merry-go-round to a rotating gravitational sphere. When we rotate the sphere, what does that do to the people trying to read their weather maps? How much rotation do they feel? The best way to envision the Earth's rotation is to see or imagine a **Foucault's pendulum**. The pendulum has to be very long so that it can keep swinging all day on a single push, so they are usually set up in an atrium 10 or 20 m high. The pendulum swings with an enchantingly long, slow period of maybe 30 seconds. The weight at the end of the wire is a

rather heavy 100 kg or thereabouts. The wire is suspended from a swivel to allow the mass on the end of the wire to rotate freely. At the base of the pendulum, a museum employee sets up a circle of little dominos, which the pendulum knocks over as it swings. Over the course of the day, the swing direction of the pendulum changes, knocking over a new little block every 10 minutes or some other crowd-pleasing rhythm. Leon Foucault installed the first of these into the Pantheon in Paris for the 1850 Paris Exposition.

The pendulum is like a rotation meter, measuring the rate of rotation that drives the Coriolis effect. Start with a Foucault pendulum on the North Pole (Figure 6-7). The Earth is spinning




Figure 6-7 The Coriolis acceleration that we feel on Earth depends on latitude. A Foucault's pendulum is a rotation detector. A pendulum set in motion at the pole would swing through 180 degrees in 12 hours, knocking down all of a circle of dominos placed around the pendulum. On the equator, the pendulum is started in a north-south direction and maintains that direction of swing as the Earth rotates. The rate of rotation, and therefore the Coriolis acceleration, is strongest at the poles. At the equator there is no apparent rotation.

relative to the fixed stars, but the swing axis of the pendulum stays stationary relative to the fixed stars. As the Earth turns under the pendulum, it knocks over the dominos. They will all be knocked over after the planet has rotated 180 degrees, which takes 12 hours.

Now move the pendulum to the equator, and start it swinging in a north-south direction. The swing axis of the pendulum cannot remain completely motionless with respect to the fixed stars because the direction “down” keeps changing as the planet rotates. But the pendulum will keep its alignment with the stars by continuing to swing in a north-south direction. The swing of the pendulum does not appear to rotate to an observer on Earth. The dominos are safe.

The atmosphere and ocean feel the earth's rotation at the poles
but not at the equator.

In the middle latitudes, the rate of rotation is between these extremes of the pole and the equator. The dominoes are completely knocked over in 12 hours at the pole, but they are never knocked down at the equator. In middle latitudes, the knock-down time is longer than 12 hours and shorter than forever. To be technical about it, the rotation rate scales as the trigonometric function sine of the latitude.

Rotations are represented in physics as arrows called **vectors**. The Earth's rotation vector would be like a flagpole, pointing not necessarily straight upward but toward the North Star, from anyplace in the northern hemisphere. The rotation that the fluids feel in the atmosphere and ocean would be represented as the height of the tip of the flagpole above the ground. On the equator, the North Star is at the horizon, and the flagpole would be lying on the ground, so its tip would have zero altitude, meaning no rotation. At the North Pole, the flagpole points straight up, achieving its full height, and the flows feel the full rotation rate of the Earth. In between, the rotation goes as the height of the tip of the pole above the ground.

Modeling the Weather

At the heart of it, fluid flow is governed by Newton's laws of motion. Because fluid has mass it has inertia, which tends to keep it moving if it is moving, or keep it stationary if it is already stationary. To change the speed or direction of the flow, motion requires a force, such as gravity or pressure differences.

One example of driving a flow by differences in pressure would be to fill up a bathtub with water and then somehow pull up on the water surface on one side, making a slope (Figure 6-8). The water will want to flow in a direction to flatten out the surface. Let us assume that the




Figure 6-8 A thought experiment in which the bottom of a bathtub is a level geopotential surface, but the water level at the surface is sloped. In this situation there would be higher pressure at the bottom of the bathtub on the deeper side than on the shallower side. This results in a pressure force that tends to push water from deeper to shallower.

bottom of the tub is flat relative to the gravity field, like the surface of a pool table, what is called a **geopotential surface**. The pressure at the bottom of the bathtub would be higher on the deep end because there is more water overhead.

The flow is driven by the **pressure gradient**, which could be calculated as

$$\text{Pressure Gradient} = \frac{\Delta P}{\Delta x} = \frac{P_{\text{bathtub floor, left}} - P_{\text{right}}}{\text{bathtub length}}$$

Our intuition is that the fluid will not stay tilted like that and will eventually flatten out. When that happens, the pressure gradient goes to zero. Because the bathtub flows more quickly than the Earth rotates, it does not feel the Earth's rotation, so the Coriolis acceleration is much smaller than the real acceleration, and the water just flows downhill.

Flows in the atmosphere and ocean persist long enough for the rotation to play a role (Figure 6-9). Flows are pushed by driving forces, such as the pressure gradient in the bathtub, or ocean flows can be driven by friction with the wind. Imagine a sudden wind pushing the water in the surface ocean. The water initially flows downwind, and as it does, the Coriolis acceleration begins to try to deflect the flow (to the right in the northern hemisphere).

Eventually, after a few rotations of the Earth, the current will reach a condition called a steady state, in which the forces acting on the water come into balance with each other. In the steady state, the current is sideways to the wind. The Coriolis force on the current is sideways to the current, pointing back upwind, balancing the wind.

A geostrophic flow balances the forces on it against each other.

The astonishing implication is that in a rotating world the fluid will eventually end up flowing completely crossways to the direction that it's pushed. This condition is called **geostrophic flow**. It is as if a marble, instead of rolling down a hill, would roll around and around the hill. It is as if the sloping surface in the bathtub would drive a flow across the bathtub. The water in the bathtub does no such thing, of course, because the flow in the tub does not have time to find the geostrophic equilibrium shown in Figure 6-9 and also because the bathtub is not infinitely wide, and water cannot flow through the walls.

Great ocean currents that persist for longer than a few days, though, do flow sideways to their driving forces. Sea level on the east side of the Gulf Stream in the North Atlantic is about 1 m higher than it is on the west side. Did I get that right? Figure out the direction that the pressure would be pushing the water, and verify that the flow is 90 degrees to the right.

Did you ever wonder why winds blow around a low-pressure cell on a weather map instead of directly from high pressure to low?

You can see geostrophic flow on a weather map as cells of high and low pressure with flow going around them (Figure 6-10). A low-pressure cell has a pressure force pointing inward all around its circumference. Ninety degrees to the right of that, the winds flow counterclockwise (in the northern hemisphere). Meteorologists call this direction of flow **cyclonic**. Flow around a high-pressure cell would be **anticyclonic**. The fluid flows in circles around the cells rather than into or out of the cells, preserving rather than wiping out the hole in the pressure of the atmosphere. This is why hurricanes can last for many days, and a huge geostrophic circulation system on Jupiter, the Great Red Spot, has probably been going for centuries.




Figure 6-9 The Coriolis acceleration affects the way that winds and currents respond to pressure forcing on a rotating planet. In the *top panel*, the fluid initially flows in the direction that the wind or pressure force is pushing it. As it starts flowing, it generates a Coriolis force directed 90 degrees to the direction of its motion. In the *middle panel*, after a while, the Coriolis force swings the fluid flow toward the right. Eventually, the fluid itself flows 90 degrees to the wind or pressure force, and the Coriolis force just balances the wind or pressure force, in the *bottom panel*. This is the steady state, where the flow stops changing and remains steady. In the southern hemisphere, the direction of the Coriolis acceleration and the steady-state flow would be reversed that shown here.

The equations that govern turbulent fluid flow cannot be solved algebraically, the way the Layer Model can, for example, but they can be solved approximately using computer models. The domain of the model, the Earth's atmosphere and ocean, is divided up into a three-dimensional grid (Figure 6-11). At each grid point, the computer keeps track of the things that drive the flow and the weather: the temperature, pressure, flow velocity, and other variables, such as water vapor concentration for the air or salinity for the ocean. At each time step, the computer keeps track of how the flow would change the temperatures and pressures and how the flow itself will change in the next time step. A typical time step for a climate model may be a few minutes.




Figure 6-10 How pressure variations and the Coriolis acceleration affect weather in the atmosphere. A region of low atmospheric pressure (*left panel*) is surrounded by pressure gradients pointing inward as air tries to flow from high to low pressure. The steady-state response to that pressure forcing is flow in a direction at 90 degrees to the right of the pressure (see Figure 6-9), resulting in counterclockwise flow around the pressure hole. The direction of flow would be the opposite around a high-pressure region (*right panel*), and both would be reversed in the southern hemisphere.




Figure 6-11 Surface wind field from a state-of-the-art climate model (a model called FOAM [Fast Ocean/Atmosphere Model], courtesy of Rob Jacob). This is a snapshot of the winds averaged over one day in the middle of January. The figure is lightly shaded to show sea-level pressure, with darker shading indicating lower pressure. One can see the wind blowing counterclockwise around low-pressure areas in the northern hemisphere, for example, just to the south of Greenland. In the southern hemisphere, winds blow clockwise around low-pressure areas, for example, in the Southern Ocean south of Australia.

Even before computers, researchers like Lewis Richardson dreamed of using human “computers” to perform the immense calculations that would be required to forecast the weather. Doing the math by hand is so slow, however, that it would never be possible to even keep up with the real weather, let alone make a forecast of the future. ENIAC (Electronic Numerical Integrator and Computer), one of the first electronic computers constructed in 1945, was used for a range of scientific computation projects, including weather forecasting.

If you would like to play with a climate computer model, you can download one called EdGCM (Educational Global Climate Modeling; see Further Reading). The model runs on a desktop computer and can simulate a few years of model time in 24 hours of computer time. It has a very nice point-and-click interface to plot maps and time series of model variables.

The physics of air and water flow, by their very nature, are challenging to simulate. This is because there are patterns in the flow that are as large as the planet and as small as a few millimeters. All the different size scales of flow are important to the overall properties of the flow and how it carries things around like heat and water vapor. Fluid turbulence is characterized by a phenomenon called a **turbulent cascade**, in which larger-scale flows break up into smaller flows. The energy in the flow passes from large scales to small, ultimately dissipating in fluid friction attributable to the **viscosity** or “goopiness” of the fluid.

Big whorls have little whorls,
Which feed on their velocity,
And little whorls have lesser whorls,
And so on to viscosity.

—Lewis Richardson

The flows in the atmosphere and ocean are ultimately forced by differences in temperature and heat flow, factors that arise on a global spatial scale. The global circulation patterns then tend to break up into smaller features as the energy trickles down the turbulent cascade. In the atmosphere, the flows break up into storms and fronts, what’s called the weather. In the ocean, large features like the Gulf Stream spawn rings and eddies a few hundred kilometers across. In atmospheric and oceanic fluid flows, there are features spanning a huge range of sizes.

Simulation of weather and climate remains one of the grand challenges in the computational sciences. This is because the mechanisms that govern fluid flow often operate at fairly small spatial scales. If we wanted a model to include everything that governs how cloud drops form and grow, for example, we would have to have grid points every few meters in the clouds. Fronts in the ocean allow fossil fuel CO₂ to mix into the ocean interior, but these can be only a few meters wide. To really resolve the mechanisms that control the climate would require an immense number of grid points in the model.

The trouble is that if we increase the number of grid points, say by a factor of 10 in each dimension, the total number of grid points goes up by a factor of $10 \cdot 10 \cdot 10$ or 1,000. That is 1,000 times more math the computer has to do per time step. To make matters worse, as the grid gets more closely spaced, the time step has to get shorter. A grid 10 times finer would require a time step about 10 times shorter. So it would take 10,000 times longer to do a model year of time. State-of-the-art computer models of climate are run at higher resolution all the time, and in general they look more realistic as resolution increases. But they are still far from the resolution they would like and will be so for the conceivable future.

Climate modelers always leave the computer room hungry for more.

Some of the physics of the real atmosphere cannot be explicitly resolved, so they must be accounted for by some clever shortcut. The real formation of cloud droplets depends on small-scale turbulent velocities and other information that the model is unable to simulate. So the model is programmed to use what information it does have, ideally in some intelligent way that might be able to capture some of the behavior and responses of the real clouds. Cloud formation may be assumed to be some simple function of the humidity of the air, for example, even though we know reality is not so simple. The term for this approach is **parameterization**. The humidity of the air is treated as a parameter that controls cloudiness. Other important parameterizations include the effects of turbulent mixing, air-sea processes such as heat transfer, and biology modeling, which is full of parameterizations. Some of these topics come up again in Chapter 11.

TAKE-HOME POINTS

- The energy budget to space of a particular location on Earth is probably out of balance, fluctuating through the daily and seasonal cycles and with the weather. This is in contrast to the Layer Model.
- The annual average energy budget for some location on Earth may not balance, either, because excess heat from the tropics is carried to high latitudes by winds and ocean currents.
- The global warming forecast requires simulating the effects of weather, which is a really difficult computational challenge.

STUDY QUESTIONS

1. How can we predict a change in climate 100 years from now if we cannot forecast a rainstorm more than 10 days out?
2. Is the Earth's energy budget to space in balance? Answer: Yes and no. Your task: Explain the answer.
3. Draw a diagram describing why geostrophic flow goes 90 degrees to the direction of the forcing.
4. Draw a diagram of a hurricane, showing the pressure gradient, the wind direction, and the Coriolis acceleration of the wind.
5. Why does the turbulent cascade make Earth's climate a difficult computational challenge?

FURTHER READING

Glieck, J. *Chaos: The Making of a New Science*. New York: Penguin Books, 1988.

EdGCM, <http://edgcm.columbia.edu/> National Aeronautics and Space (NASA) scientists have created a climate GCM model that you can run easily on Windows or on the Macintosh. There

are graphical menus for setting up the model run: greenhouse gas concentrations, solar intensity, and so on. There are also graphical tools to make maps and time series of the model output. I can run a few decades of model time overnight on my relatively old laptop.

Exercises

1. **Earth's Orbit and Seasons.** Answer this question using an online model of the intensity of sunlight as a function of latitude and season at <http://understandingtheforecast.org/orbit.html> (Orbit/Snapshot on the Forecast/Projects Web page). The model calculates the distribution of solar heating with latitude and season depending on the orbital parameters of the Earth. Enter a year AD and push "calculate." The eccentricity is the extent to which the orbit is out-of-round; an eccentricity of 0 would be a fully circular orbit. Obliquity is the tilt of the Earth's axis of rotation relative to the plane of the Earth's orbit. The third number, the longitude of the vernal equinox, determines the location on the orbit (the date of the year) where the northern hemisphere is tilted closest to the sun. Using the present-day orbital configuration, reproduce Figure 6-4. Which of the three orbital parameters is responsible for the seasons? Find an orbital configuration that has no seasons, or much reduced seasons, and explain why your scheme works.
2. **Heat Transport.** Answer these questions using an online full-spectrum radiation model at http://understandingtheforecast.org/Projects/full_spectrum.html.
 - a) The incoming solar radiation at the equator, averaged over the daily cycle, is about 420 W/m^2 . The default albedo is set to be 30% in the model; change this to a more realistic value for the tropics of 10%. What would the temperature be at the equator if there were no heat transport on Earth? The temperature in the tropics of the real world is about 30°C , a lower answer than you probably got because of heat transport from the tropics to the pole, represented in the field labeled "TOA (for Top of Atmosphere) radiative imbalance," where a negative number indicates heat export. How much heat export is required to get a reasonable temperature for the equator? What fraction of the solar influx is this?
 - b) Repeat the same calculation for high latitudes. Estimate the annual average heat influx at 60 degrees latitude by running the orbital model from Project 1. Just eyeball the fluxes through the year to guess at what the average would be. Now plug this into the full-spectrum light model to see how cold it would be up there if there were no heat transport. If there were no transport and also no storage, how cold would it be in winter?

Complexity in the Earth system arises from the way pieces of it interact with each other

Positive and Negative Feedbacks

A **feedback** is a loop of cause and effect (Figure 7-1). At the center of a feedback is a **state variable**. The state variable in many of the climate feedback loops in this book is the average temperature of the Earth. To see a feedback in action, drive the temperature a bit by changing some external factor, like the intensity of the sun. A **positive feedback** makes the temperature change larger than it would have been without the feedback, amplifying the temperature change. A **negative feedback** counteracts some of the external forcing, tending to stabilize the state variable.

Stefan-Boltzmann Feedback

The energy budget of the Earth, balancing infrared (IR) against sunlight energy fluxes, is controlled by a negative feedback mechanism (Figure 7-1a). The outgoing energy flux depends on temperature according to the Stefan-Boltzmann rule, $\epsilon\sigma T^4$. If you were to suddenly deposit a whole lot of extra energy on the planet, say if there were a giant asteroid impact that raised the temperature of the surface, the feedback would pull the temperature back down. This is a negative feedback because the feedback pulls in the opposite sense from the perturbation.

A negative feedback is a stabilizer. The drain in the kitchen sink analogy has a negative feedback to the water level.

Ice Albedo Feedback

There is a positive feedback in the climate system known as the **ice albedo feedback**, which operates on the state variable of temperature (Figure 7-1b). An input perturbation, such as a rise in greenhouse gases, drives temperature up a bit. Ice melts, reducing the albedo (the reflectivity of the Earth: Chapter 3), allowing the dark ground to absorb more sunlight than the ice would have, warming things up a bit more. This is called a *positive feedback* because the directions of the input perturbation and the feedback loop agree with each other. A positive feedback can work in the opposite direction as well, taking a cold time and make it even colder. A positive feedback loop amplifies an excursion in either direction.

A positive feedback is an amplifier.




Figure 7-1 Feedback diagrams. (a) An example of a negative feedback resulting from the Stefan-Boltzmann infrared energy flux $\epsilon\sigma T^4$. (b) An example of a positive feedback. Some external perturbation increases the temperature. The increase in temperature causes ice to melt, allowing the land to absorb more of the incoming solar radiation (by decreasing the albedo). The melting ice drives temperature up further. (c) The negative feedback that controls atmospheric water vapor, given an atmospheric temperature. (d) The water-vapor feedback, which significantly amplifies the temperature effect of rising CO_2 in our climate.

The ice albedo feedback is most important in the high latitudes because that is where the ice is. The warming response of climate models to extra CO_2 is two to four times stronger in the high latitudes than on the global average. This is also observed in the warming pattern of the last few decades, with a caveat. Warming is stronger in the Arctic than the global average, driven in part because the sea ice in the Arctic is melting rapidly. At the other pole, the interior of Antarctica is actually cooling, which may be a special case, the result of the Antarctic ozone hole (see Chapter 9) and its impact on a spinning column of air called the polar vortex. Sea ice around Antarctica is also holding relatively steady, not melting like in the Arctic. But the Antarctic Peninsula, a little farther north, has warmed more than any place else on Earth, and this is where most of the exploding ice shelves are from, like the Larsen B pictured in Chapter 12.

The presence or absence of sea ice can also have a huge impact on local temperatures because air temperatures over open water are moderated by the heat sink of the water column. As the air draws heat from the water, turbulent mixing and convection in the water will quickly refresh the surface water from below. As a result, air over water does not get much colder than freezing, 0°C , but air over ice or land can be much colder. If a landmass is frozen into a vast expanse of sea ice, it will have a climate more typical of the interior of a continent. This sea-ice temperature feedback helps explain the abrupt climates recorded in the Greenland ice cores (Chapter 12).

Water-Vapor Feedbacks

Water vapor is responsible for more greenhouse heat trapping on Earth than CO_2 is, and yet global warming seems to be about CO_2 , not water vapor. No one seems concerned about global warming from running a lawn sprinkler. Why not? The answer is that there is a negative feedback loop that controls the amount of water vapor in the atmosphere at any given temperature, having to do with rainfall and evaporation (Figure 7-1c). If there is too much water vapor in the air, it will condense out as rain. In the other direction, if the air is extremely dry, any available liquid water will tend to evaporate into it. The lawn sprinkler is not going to lead to global warming because any water forced to evaporate into the atmosphere today will just rain out next week.

But water vapor is also involved in a positive feedback loop acting on global temperature (Figure 7-1d). Because warmer air holds more water vapor than cooler air, warming allows more water to evaporate before it rains (see Figure 5-8). Water vapor is a greenhouse gas, so its increase tends to warm the planet still further. The water vapor feedback is a positive feedback, powerful enough to more or less double the climate impact of rising CO_2 concentration. If it were not for the water-vapor feedback, Earth's climate would be much less sensitive to CO_2 , and maybe we would not be worrying so much about global warming.

Water vapor traps more than CO_2 does, but it uses its power to amplify climate changes driven by CO_2 .

It is possible for the water-vapor feedback to feed into itself, running around the loop of cause and effect in what is known as a **runaway greenhouse effect**. It happened on Venus, but we are in no immediate danger of it happening on Earth. A runaway greenhouse effect stops if the vapor concentration in the air reaches saturation with liquid water or ice, so that any further evaporation would just lead to rainfall or snow.

Figure 7-2 is called a **phase diagram**, and it shows what phases of water you will find as a function of temperature and pressure. The **phases** of water are solid, liquid, and vapor.




Figure 7-2 A phase diagram for water demonstrating that the water-vapor feedback on Earth and Mars is limited, whereas Venus is free to have a runaway greenhouse effect.

At high temperature and low pressure, you expect to find only vapor. Vapor likes it hot. Cool it down and increase the pressure, and liquid will condense, and when it gets colder, ice will form. At the triple point, a specific pressure and temperature combination, you get all three phases coexisting together, a glass of boiling ice water. Of course, you can throw ice cubes in a pot of boiling water on your stove anytime you like, but the ice will quickly melt because it is not at equilibrium. The phase diagram in Figure 7-2 shows equilibrium states.

To understand the water-vapor feedback, imagine what happens if you suddenly introduce water to a planet that initially had none. This starting position on Figure 7-2 is all the way at the bottom of the figure, where the pressure of water vapor in the air is low. From here, what will happen is that water will evaporate and the vapor pressure in the atmosphere will rise. Water vapor is a greenhouse gas, so as the water-vapor content of the atmosphere increases, the temperature warms. Therefore, as the condition of the planet moves upward on the figure, to higher water-vapor pressure, it veers to the right, which is higher temperature.

The amplifying loop of cause and effect stops if the vapor pressure and temperature hit the liquid or ice stability fields.

The middle of the three curves on Figure 7-2, labeled Earth, moves up and to the right until it intersects the stability field of water. At this point, the atmosphere is holding as much water vapor as it can. Any further evaporation just makes it rain. The curve on the left on Figure 7-2 represents something like the situation on Mars. Here the water-vapor feedback path intersects the stability field of ice.

The only planet on this diagram with a runaway greenhouse effect is Venus, the curve on the right. Water on Venus would evaporate, increasing the temperature, as before. The difference is that the path that Venus follows never intersects the stability fields of either liquid or solid water. Planetary scientists presume that Venus originally had about as much water as Earth but that the high solar heat flux associated with orbiting so close to the sun drove that water to evaporate into the atmosphere, rather than condense into oceans as it has on Earth.

Evaporation of all of a planet's water is a one-way street because, if water vapor reaches the upper atmosphere, its chemical bonds will get blown apart by the intense ultraviolet light. The hydrogen atoms, once they are disconnected from oxygen, are small enough and light enough that they are able to escape to space. The water is lost for good. This is the presumed fate of Venus's water.

A runaway greenhouse effect means the end of a planet's water.

Earth has retained its water as a result of the temperature structure of the atmosphere. It gets colder aloft until you reach the tropopause, the boundary between the troposphere and stratosphere (Figure 5-2). The tropopause acts as a cold trap for water vapor, making sure that most of it rains or snows out before the air rises too close to space. The oceans, viewed from space, look heartbreakingly vulnerable, but they are protected, and apparently have been for billions of years, by a thin layer of cold air. Marvelous!

Theoretically, if the sun were to get hotter, or if CO₂ concentrations were high enough, the Earth could move to the right on Figure 7-2, sufficiently far that it could escape the liquid water stability field and hence run away. But don't worry; there is not enough fossil-fuel carbon on Earth to do this. The sun is heating up over geologic time, but only very gradually, and the runaway greenhouse effect from this is nothing to worry about for billions of years.

Interestingly, if the equator were isolated from the poles, blocked from any heat transport and forced to balance its energy fluxes using outgoing IR only, the tropics would be a runaway greenhouse.

Earth's climate uses the high latitudes as cooling fins to avoid the runaway greenhouse effect.

The role that the water-vapor feedback plays in the global warming forecast can be drawn on this same diagram, although global warming climate change is tiny compared with Venus's climate, so I have drawn it horribly exaggerated so it can be seen. An initial temperature perturbation, say from rising CO₂, moves the water-vapor feedback path a bit toward the right, toward Venus. The new path ends as it hits the liquid saturation line, the way it did for the original path. The temperature difference between the two raining, moist Earths (where they intersect the liquid water stability field) is greater than the temperature difference between the two dry Earths (at the bottom of the plot). The water vapor feedback amplifies the temperature change that you would get from increasing CO₂ on a dry planet.

On the real Earth, mercifully, the relative humidity is not everywhere 100%. One mechanism controlling the distribution of atmospheric water vapor is the Hadley circulation (Figure 7-3). Warm air at the equator rises convectively. Water condenses as the air rises and cools. The rising column of air generally has a lot of water vapor in it. Then the air spreads out at high altitude and begins to subside in the subtropics, about 30 degrees latitude north and south. This is air that has been through the "wringer," the cold tropopause, and there is not much water vapor left in it. The great deserts of the world are located under these dry air subsidence regions. The air flows back equator-ward along the surface, picking up water vapor as it goes. In general, the humidity of the atmosphere depends on the patterns of circulation and access to evaporating water.

Globally, the average humidity of surface air is about 80%. The real Earth, if it were plotted on Figure 7-2, might allow the humidity to rise to 80% of the way to the liquid-phase boundary




Figure 7-3 The Hadley circulation and its effect on the humidity of the air in the troposphere.

rather than making it all the way to 100%. But even this is too simple because water vapor is most important as a greenhouse gas at high altitude, where the air is cold. Relative humidity in the coldest air can be lower than at the surface, as dry as 10% relative humidity in the subsiding areas.

The strength of the water vapor feedback depends on what happens in the upper atmosphere, where water vapor is most important as a greenhouse gas.

Climate models tend to predict that the relative humidity of the atmosphere should stay about the same as the temperature rises. Because the saturation vapor pressure increases with temperature, the absolute humidity (number of molecules of H_2O per molecules of air) rises with temperature. So the water vapor feedback still operates, even though the air is not everywhere 100% saturated.

It is possible that the models are wrong, and the absolute humidity could change more or less than this. One could imagine changing the relative humidity by changing the strength of the Hadley circulation, or the rate of turbulent mixing of air by storms. Is it possible that the water-vapor feedback in the real world could be weaker than in the models? If so, it would make a big difference in the amount of climate change to expect. An ultimate answer to that question comes from paleoclimate and CO_2 reconstructions, discussed in Chapter 12.

Clouds

There are three main types of clouds (Figure 7-4). High-altitude clouds are called **cirrus clouds**. These are thin and wispy, often barely noticeable in the sky, made of ice crystals. There are also two major types of low-altitude clouds. **Cumulus clouds** are towers, the result of focused upward blasts of convection. Thunderstorms come from cumulus clouds. **Stratus clouds** are layered, formed by broad diffuse upward motion spread out over large geographical areas.

Clouds interfere with both the visible light and the IR energy fluxes in the Earth's radiation balance. In the IR, clouds act as blackbodies (Chapter 3), warming the planet. The effect on the visible-light energy budget is to reflect light to space, cooling the planet. The overall impact of a cloud depends on which of these two effects is stronger, which in turn depends on what type of cloud it is.

Clouds are pretty good blackbodies, absorbing and potentially emitting just about any frequency of IR. This makes them different from greenhouse gases, which are choosy about what frequencies of IR radiation they interact with. Clouds are blackbodies because they are made of liquid droplets or ice crystal solids. In general, liquids and solids at the Earth's surface tend to be black in the IR.

The effect of a cloud on the IR energy budget depends on the temperature at the top of the cloud, where IR radiation heading upward escapes the cloud and heads to space. High clouds emit IR radiation at cold temperatures, blocking the bright IR from the warm ground. Clouds




Figure 7-4 Schematic of the three main types of clouds.

with low tops have a smaller effect on the outgoing IR light. Table 7.1 records the conclusions that the effect of clouds on IR light is to warm the planet and that high clouds warm the Earth more than low clouds do.

The infrared effect of clouds (1) warms the Earth (2) depending on the altitude (temperature) of the cloud tops.

The effect of clouds in the visible-light energy budget is to send solar energy back to space, increasing Earth's albedo and cooling the planet. Cloud particles (droplets and ice crystals) can also simply absorb the sunlight, transforming its energy into heat in the particle. But absorbing the heat in the atmosphere instead of absorbing it at the ground has a smaller impact on Earth's energy budget than reflecting light back to space without ever absorbing it.

Light can be scattered when the electric field of the incoming visible light causes the electric field in the particle to oscillate. The oscillating electric field in the particle then emits visible light, the same frequency as the light that came in. The difference to Earth's energy budget between absorbed and scattered sunlight is that when light is scattered back to space, its energy is never converted to heat, so it never enters into the planet's heat energy budget.

The fraction of incoming sunlight that is scattered versus that absorbed by a cloud varies a lot between different types of clouds. You cannot see through low clouds, meaning that they are optically thick, capturing or scattering most of the visible light that tries to get through. Cirrus clouds contain 10 or 100 times less water per volume than lower altitude clouds typically hold.

Table 7.1

	High Clouds (Cirrus)	Low Clouds (Stratus and Cumulus)
Infrared effect (warming)	Strong warming influence because high-altitude cloud top	Weaker warming influence because cloud tops are lower
Visible light effect (cooling)	Weak cooling influence because they are optically thin	Stronger cooling influence because they are optically thicker
Overall effect	Warming	Cooling

They are so thin you can often see blue sky right through them. Cirrus clouds therefore have a weaker cooling impact on the visible-light budget than high clouds do (Table 7.1).

Putting together the IR and visible parts of the energy budget, high clouds warm the Earth and low clouds cool it. The greenhouse effect wins for high clouds, and the albedo effect wins for low clouds. Overall, the cooling effect of the low clouds dominate, leading to a net cooling of Earth's climate.

When climate models are subjected to higher atmospheric CO₂, their clouds tend to respond by amplifying the warming from the CO₂ in a positive feedback. No models predict a negative feedback from clouds, but the strength of the feedback is very uncertain, and is the main source of uncertainty in model calculation of the climate sensitivity of the Earth.

High clouds warm, low clouds cool.

The scattering efficiency of clouds varies a lot with meteorological conditions, and it is also affected by human pollution. The important factor is the cloud droplet size. Scattering would be most efficient if the droplets were about the same size as the wavelength of the light. But cloud drops are almost all much larger than the wavelength of visible light, so the smaller the drop, the better for scattering. Rain clouds look dark because their large droplets do not scatter light very well. Changing the cloud droplet size from 8 to 10 μm in the online National Center for Atmospheric Research, or NCAR, full-spectrum radiation code (Exercise 7-3) changes the energy budget by several W/m^2 , a radiative forcing comparable to that of doubling CO₂.

Different types of cloud droplets scatter light in different directions. Most of the light scattered by spherical liquid drops continues in a more or less forward direction, which, if the incoming direction is downward, would mean to continue downward toward the Earth. Ice crystals are better at actually reversing the direction of light, sending it back up into space.

The size of cloud droplets can be affected by **cloud condensation nuclei**, “seeds” that help droplets form. There are natural cloud condensation nuclei from sea salt, dust, pollen, smoke, and sulfur compounds emitted by phytoplankton. The strongest human footprint in the clouds comes from coal-fired power plants releasing sulfur. Sulfur is emitted in flue gas as SO₂, and it oxidizes in a week or so to sulfuric acid, H₂SO₄, which condenses into small droplets less than 1 μm in size called **sulfate aerosols**. These are so small that they scatter a lot of the light that encounters them.

The aerosols act as cloud condensation nuclei because the strong acid tends to pull water out of the air like those silica gel packets that ship with electronics equipment. Humans also produce cloud condensation nuclei from internal combustion engines, especially on ships, and from forest fires and cooking fires. Most of the near-surface atmosphere has enough natural condensation nuclei that the issue is not whether to form a droplet or not when the air gets supersaturated. But adding more condensation nuclei will encourage the water to form a greater number of smaller-sized droplets. The climate impact of changing the cloud droplet size arises because smaller droplets scatter light more efficiently than larger droplets do. This is called the **sulfate aerosol indirect effect**. The climate impact of indirect aerosol effect is very uncertain, but it is probably even larger than the direct effect of scattering light by the aerosols themselves (Figure 10-2 in Chapter 10).

Clouds that form in dirty air tend to be better scatterers with a higher albedo, cooling the planet.

Contrails are another example of cloud seeding by humans. Contrails form when a jet airplane passes through clean air that is somewhat supersaturated with water vapor. The exhaust contains cloud condensation nuclei that quickly absorb water from the air to form droplets.

The contrail particles spread out and eventually become indistinguishable from natural cirrus particles. For this reason, it is difficult to know how much impact aircraft have on the cloudiness of the upper atmosphere. Contrails tend to warm the planet, just like natural high clouds. But current thinking is that the climate forcing of clouds induced by airplanes is pretty small, about 1% of the climate forcing from greenhouse gases. Contrails and sulfate aerosols both differ from greenhouse gases in that they would go away very quickly if production stopped, whereas the greenhouse gases, especially CO₂, persist.

Clouds are the largest source of uncertainty in climate models. Most of the differences in climate sensitivities between different models are due to the different ways they deal with clouds. The processes of cloud droplet formation and growth are controlled by details of the windy turbulence in the cloud, which cannot be explicitly modeled in a global model because the computational grid of the model is too coarse to represent all of the gusts of wind that shape the cloud (see Chapter 6).

So, even with the astonishing explosion of computer power available, modeling clouds globally from first principles is pretty much out of the question. Instead, clouds must be parameterized, meaning approximated, at a higher level, as functions of humidity, temperature, or whatever other variables seem to be important. Cloud models are tested by comparing the distributions of clouds today with weather data. There is a wide range of atmospheric conditions around the globe, so this is a strenuous test, but it could be that things could change in a changing climate in ways that are not anticipated. For reasons such as this, it is important to validate the models by comparing their behavior to historical changes in Earth's temperature over the past few decades or through reconstructed climate changes of the prehistoric past (Chapter 11).

Clouds are loose cannons in climate models.

Most of the uncertainty in the climate sensitivity of the Earth (temperature change that would result from a doubling of atmospheric CO₂ concentration) arises because of a positive climate feedback involving clouds, amplifying the effect of the CO₂ change. If this feedback is weak in a climate model, the climate sensitivity is relatively low, maybe 2°C. If it is stronger, the climate sensitivity is higher. However, there are no models in the Intergovernmental Panel on Climate Change, or IPCC, intercomparison that exhibit a negative feedback from clouds.

Ocean Currents

The oceans interact with climate in many ways. One example is a periodic flip-flop between two states of the ocean called **el Niño** and **la Niña** (Figure 7-5). On the atmospheric side, there is a corresponding cycle in the atmospheric pressure difference across the Pacific, called the **Southern oscillation**. Emphasizing that this is a coupled atmosphere/ocean phenomenon, the el Niño cycle has been abbreviated **ENSO**, which stands for el Niño-Southern oscillation. The state of the ENSO cycle affects patterns of climate all around the world.

In general, the world dries out during el Niño.

In the la Niña state, there is cold water at the sea surface in the eastern part of the Pacific Ocean near Peru. The cold water contains nutrients that allow phytoplankton to grow, feeding the very productive Peruvian anchovy fishery. The contrast in sea-surface temperatures between the eastern and western equatorial Pacific drives a wind along the equator that blows from east to west. The warm water, floating on the colder water with a boundary call the **thermocline**, piles up in the West. The tilted thermocline keeps the surface waters cold near Peru, which drives winds, which keep the thermocline tilted and the cold water at the surface.




Figure 7-5 The configuration of the atmosphere and ocean along the equator in the Pacific Ocean during el Niño and la Niña phases of the el Niño climate oscillation. *Top panel:* During la Niña, the boundary between warm water and cold water in the ocean slopes up from west to east, exposing cold water to the atmosphere in the east near Peru. In addition to stimulating the Peruvian fisheries, the cold water drives a strong wind from East to West that tends to sustain the slope in the ocean temperature boundary. In the el Niño phase of the oscillation (*bottom panel*), the slope in the ocean temperature boundary and the winds in the atmosphere collapse, as do the fisheries in Peru. The coupled atmosphere/ocean system flips back and forth between these two states every 4 to 7 years or so, affecting weather patterns around the rest of the world.

The cold water near Peru gets covered with warm water during el Niño. Fishermen were the first to notice the el Niño cycle because the warm waters were less fertile and the fisheries would collapse that year. The winds along the equator falter because the sea-surface temperature difference that drove them has disappeared.

Although both of these states seem to be self-stabilizing (the winds and thermocline tilt arrayed in a positive feedback), the atmosphere/ocean system flips back and forth between the two climate states, about one cycle every 4 to 7 years.

There is a possibility that with global warming the Pacific may tend to favor the el Niño state, the climate impacts of which could be considered a feedback. However, the extent and strength of this potential feedback are impossible to forecast reliably.

A warmer world might have a permanent el Niño. Or maybe not.

Another potential feedback from the ocean to climate is called the **meridional overturning circulation** in the North Atlantic Ocean. Warm water is carried to the North Atlantic in the Gulf Stream. As the water cools, its density increases and it sinks into the abyss, making room at the surface for more warm water to carry more heat from the tropics. The Pacific Ocean has a surface current analogous to the Gulf Stream, called the Kuroshio Current, but surface waters in the Pacific are not salty enough to get dense enough to sink, so there is no deep meridional overturning circulation in the Pacific.

Reconstructions of the climate and oceanography of the North Atlantic region from the last ice age point to an instability in the meridional overturning circulation. The climate records from ice cores in Greenland show huge temperature swings called **Dansgaard-Oeschger** events, changes of 10°C within just a few years. Oceanographic reconstructions from sediment cores show that the depth meridional overturning circulation in the North Atlantic changes in synchrony with these temperature swings. One large climate shift, called the **8.2k event**, 8,200 years ago, has been correlated with the catastrophic release of a Great Lake's worth of fresh water into the North Atlantic, as an ice dam holding the lake finally let water past.

Sudden changes in the overturning circulation in the North Atlantic have driven abrupt global climate changes in the past. In the future this circulation might be impacted by warming or by melting Greenland's ice.

The meridional overturning circulation in the North Atlantic may be sensitive to future changes in climate and fresh water discharge. If the Greenland ice sheet were to melt within a century (faster than expected), it would release enough fresh water to slow the overturning circulation. Climate models also predict a slowdown in the circulation with very high CO₂ levels, say four times higher than today. If the overturning circulation slows, it could lead to cooling in the high northern latitudes, driven by the loss of heat coming up from the tropics in the poleward surface current.

The climate models generally find that by the time the overturning circulation slows down in a model run, the warming from the CO₂ is much stronger than the local cooling from the collapse of the circulation. This makes the “Day After Tomorrow” scenario, of a sudden collapse into a new ice age, seem unlikely. However, climate models also tend to underpredict the severity and spatial footprint of past abrupt climate changes driven by instability of the meridional overturning circulation. So the impact of a shutdown in the future could also be stronger than the models predict. There is more on this topic in Chapter 12.

Terrestrial Biosphere Feedbacks

The terrestrial biosphere has the potential to feed back to climate if changes in vegetation alter the albedo of the land surface. Expansion of the northern forests into areas previously covered by tundra, for example, could make the land surface darker, tending to warm things up further.

The land surface affects the albedo, and it stores carbon.

Trees also impact the hydrological cycle by extracting ground water and evaporating it from their leaves, a process called **transpiration**. Rain forests in particular are thought to perpetuate their own existence by extracting ground water and evaporating it to the air, recycling it to rain again, rather than allowing it to escape in river flow to the ocean. In a drought, vegetation dies

and soils dry out, losing the ability to retain water, thereby perpetuating the water shortage in a positive feedback. It takes an extra boost of water supply to get out of a drought, overcoming this vegetation feedback.

Carbon-Cycle Feedbacks

There are many potential feedbacks between climate and the carbon cycle, the subject of the next three chapters. Warming the atmosphere increases the temperature of the surface ocean, which tends to drive CO_2 to degas from the ocean to the atmosphere. The ultimate trigger of the glacial-interglacial cycles is thought to be variations in the Earth's orbit, but the climate swings would have been much smaller if the carbon cycle had not somehow shifted, pulling CO_2 out of the atmosphere as the ice sheets started to grow in a positive climate feedback.

The land surface stores a lot of carbon. If the Amazon or Indonesian rain forests were to burn, for example, they could release a significant amount of carbon as CO_2 to the atmosphere. Other sources of carbon that might feed back to climate change include peats, organic carbon in permafrost soils, and methane hydrates in the ocean. The methane cycle could also feed back to climate if there are climate-driven changes in wetlands, for example, or changes in the radical oxidation chemistry of the atmosphere (Chapter 10).

Feedbacks in the Paleoclimate Record

In general, first-attempt climate models seem to have an inherent tendency to underpredict the extremes of climate variation in the real climate, as inferred from climate records such as ice cores. I believe this is due to the difficulty of capturing feedbacks in climate models. Feedbacks often couple together very different parts of the climate system, requiring cloud physicists to work together with an oceanographers, for example. It takes creativity to think of some of these relationships! The message is that the future will certainly surprise us in some ways.

TAKE-HOME POINTS

- Positive feedbacks act as amplifiers of variability, whereas negative feedbacks act as stabilizers.
- The water-vapor feedback doubles or triples the expected warming owing to rising CO_2 concentrations.
- The ice albedo feedback amplifies the warming in high latitudes by a factor of three or four.
- Clouds have a potentially huge impact on climate. Clouds are expected to exert an amplifying feedback to climate warming, although the strength of this feedback is uncertain. Clouds are the largest source of uncertainty in model estimates of the climate sensitivity.

STUDY QUESTIONS

1. Explain why a stabilizing feedback is called “negative.”
2. Describe and characterize as positive, negative, or unknown feedbacks associated with ice, water vapor, clouds, ocean circulation, and the carbon cycle.
3. Is Earth's climate vulnerable to a runaway greenhouse effect? Why or why not?
4. What factors determine the climate impact of a cloud?

FURTHER READING

Kump, L. R., J. F. Kasting, P. G. Crane. *The Earth System*. Saddle River, N.J.: Prentice Hall, 1999.

Walker, J. C. G.. *Numerical Adventures with Geochemical Models*. Oxford, U.K.: Oxford Press, 1991.

Exercises

1. **Compare two radiation codes.** Global climate models handle energy carried by light, both visible and IR, in a much simpler way than the line-by-line calculations that Modtran does for IR. Modtran would be too slow to run at every grid point of a global model. Find the full-spectrum radiation code at http://understandingtheforecast/Projects/full_spectrum.html. The page can run either of two radiation schemes, called CCM and Chou.
 - a) Run the model twice, one with default conditions and the other with doubled CO₂ concentration, to determine the climate sensitivity, ΔT_{2x} , for each of the two models.
 - b) Run one of the models comparing the temperature effect of doubling CO₂ from 100 to 200 ppm. Do you get the same ΔT_{2x} as you got for 350 to 700 ppm?
 - c) The model includes the water vapor feedback automatically, but you can turn this off by zeroing the relative humidity. What is the ΔT_{2x} without the water vapor feedback?
2. **Clouds and infrared.** Use the online IR radiation model http://understandingtheforecast/Projects/infrared_spectrum.html.
 - a) Add some altostratus clouds, then stratus clouds, using the cloud pull-down menu. Describe the effect of the clouds on the outgoing IR spectra between the three cases. Draw a Layer Model-type cartoon showing the ground, the cloud, the CO₂, and two flavors of IR light at 700 and 900 cycles per centimeter (the CO₂ bend and the atmospheric window). Which type of cloud has the greatest impact on the outgoing IR flux and why?
 - b) Set the sensor altitude to 0 km, and choose the “Looking Up” option. Do this first with no clouds. You see IR light coming down from the sky. Where does it come from? How can you turn it off? Now turn on some clouds and describe how this affects the IR spectrum. Why is it warmer at night when there are clouds?
3. **Clouds and visible light.** Use the NCAR model to calculate the radiative forcing, in W/m², from high and low clouds.
 - a) First determine the equilibrium ground temperature with no clouds. Then add 100% high cloud cover. The temperature will change severely. Change it back to the original temperature by dialing up or down the top of atmosphere (TOA) radiative imbalance. This tells you how much the clouds change the balance of energy fluxes.
 - b) Do the same thing for low clouds. Clouds affect both the IR and the visible light energy fluxes: which effect is strongest from each type of cloud?
 - c) With the low clouds still at 100% and the TOA radiation imbalance still set to keep a temperature near the default 16°C, change the drop size from 10 to 8 μm, and find the radiative forcing again, by dialing up or down the TOA to bring the temperature back as before. Now turn clouds off and double atmospheric CO₂, and find the radiative forcing of that. How do the radiative effects of cloud droplet size compare with the effect of doubling CO₂? You can see how important it is to get the cloud droplet size right in a climate model!

Part II

The Carbon Cycle

It's alive!

The Chemistry of Carbon

Carbon has rich, spectacular, wondrous chemistry. There are more scientists studying the chemistry of carbon than any other element, I am sure, if you count all the doctors and biologists. The nearest relative of carbon is the element silicon, which is located directly underneath carbon on the periodic table. Silicon is one of the more abundant elements on Earth, and the chemistry of the rocks in the Earth's mantle and crust is controlled to a large extent by the whims of silicon. Silicon in the environment finds its way into hundreds of different crystal forms, including the minerals in rocks and clays. However, silicon chemistry is regimented, flat—one could almost say lifeless—compared with the chemistry of carbon.

Carbon chemistry is kept highly organized within living things, but after life is finished with it, carbon left in soils, sediments, and rocks forms itself into a disordered, indescribable goo called **humic acids** or **kerogen**. Silicon forms hundreds of crystals, but the number of possible configurations of carbon is essentially infinite. I scoff at the cliché from *Star Trek*, “carbon-based life forms”, spoken as though this were some sort of novel idea. Certainly I am close-minded and lack vision, but I would bet against finding life based primarily on any element other than carbon.

Planetary scientists presume that the three terrestrial planets—Venus, Earth, and Mars—all received about the same amounts of carbon when they formed. The carbon on Venus is mostly found in the atmosphere, and there is so much carbon that the pressure on the surface of Venus from CO₂ is 70 times higher than the total atmospheric pressure on Earth. The carbon on Earth is mostly found in limestone sedimentary rocks, with only about 0.039% of 1 atmosphere of pressure in the air (390 microatmospheres). The amount of CO₂ in the atmosphere of Venus is 180,000 times higher than the amount in Earth's atmosphere. Clearly, a planet has a wide range of options about where to put its carbon.

Carbon on Earth is the backbone for constructing the machinery of life. In addition to serving as the framework and scaffolding for that machinery, carbon chemistry also provides the means of storing energy. Photosynthesis converts the energy of sunlight into high-energy carbon biomolecules. This energy might be used by the plant, or perhaps an animal eats the plant and uses its energy for itself. Over geologic time scales, burial of biological carbon stores energy on a planetary scale, “charging up” the biosphere with a store of biomolecules and atmospheric oxygen like a gigantic battery. When we extract energy from fossil fuels, we harvest some of the ancient energy stores of the biosphere. In the process we rearrange the distribution of carbon among its various reservoirs on Earth.

To understand how all of this works, you need to know about a fundamental property of atoms in molecules, called **oxidation state**. An example of oxidation state that you are familiar with is that iron oxidizes when it rusts. Three simple carbon molecules demonstrate the spectrum of reduction-oxidation (abbreviated **redox**) chemistry for carbon.

	Oxidized		Reduced
Simplest example	CO ₂	CH ₂ O	CH ₄
Carbon oxidation state	+4	0	−4
General category	Inorganic carbon	Carbohydrates	Hydrocarbons

The first is carbon dioxide (CO₂), a greenhouse gas. The last is methane, also a greenhouse gas. CO₂ is the **oxidized** form of carbon in this spectrum, and methane is the **reduced** form. You might guess that the term *oxidation state* implies that we must be talking about oxygen in some way, and in fact the oxidized form of carbon has oxygens on it.

Oxidation and reduction has to do with bookkeeping of electrons.

In more general terms, however, oxidation state is a measure of the surplus or deficit of electrons around the carbon. Oxygen is a very greedy element for the two electrons that it requires to find its most stable electronic configuration. Usually oxygen in a molecule is credited with stealing two electrons from its bond partner, although if its bond partner is another oxygen molecule, they must share the electrons equally. The carbon in CO₂ has two oxygen partners, so it has given up four electrons, two to each oxygen atom. Because each electron has a charge of −1, a deficit of four electrons leaves the carbon in CO₂ with an **oxidation state** of +4.


Hydrogen, in contrast, is relatively generous with its single electron. The four hydrogens in methane each donate one electron to carbon, so carbon ends up with an oxidation state of −4. The oxidation state of CH₂O, the middle compound in the spectrum, is zero, because carbon gains two electrons from the hydrogens but donates two to the oxygen.

Many other carbon compounds have these oxidation states, and it makes sense to group them together into families in this way. Most of the carbon on Earth is oxidized carbon, also called **inorganic carbon**, in CaCO₃ (limestone) rocks or dissolved in the oceans. Inorganic carbon has a fascinating chemistry (spoken as one who has devoted his life to understanding it), but it is nowhere near complex enough to create the machinery of life.

At the reduced end of the spectrum, **hydrocarbons** include oil as well as natural gas. Some biomolecules are hydrocarbons, such as fats. Again, amazing stuff, but life could not be constructed from fats alone. Life is built from carbon largely in the intermediate oxidation state, called **carbohydrates**. Sugars, for example, have formulas that are multiples of CH₂O (a chemical called formaldehyde), like glucose which is (CH₂O)₆.

Hydrocarbons and carbohydrates are together referred to as **organic carbon**. Here is a case where scientific and common usages of a word are completely different. To a chemist, compounds made of reduced carbon, such as dichlorodiphenyltrichloroethane (DDT), dioxin, and polychlorinated biphenyls (PCBs), are considered organic compounds. Historically, reduced carbon was called organic because it was thought that only life could make it. Today, in common usage, “organic” produce at the grocery store is supposed to be free of toxic manmade compounds such as these.

The most stable form of carbon in the chemical environment of the Earth’s surface, in the presence of oxygen, is the oxidized form. The carbon pools on Venus and Mars are also both nearly entirely oxidized. Reduced carbon at the surface of the Earth comes from **photosynthesis**, which harvests energy from sunlight to make organic carbon



The energy from light is absolutely required; photosynthesis cannot happen in the dark, say in the deep sea or buried in soils. The only forms of life on Earth that do not owe their chemical energy to sunlight are communities living near hot-water springs in the deep ocean, called **deep sea vents**. These vent communities ultimately get their energy by catalyzing chemical reactions between the vent fluid and the seawater. Life may have gotten its start in this way, and it may exist in this way on other worlds, but on Earth photosynthesis is by far the dominant source of energy for life.

Carbon at the Earth's surface wants to be oxidized. It takes energy to make reduced carbon.

The backward direction of the chemical reaction is called **respiration**, the only energy source for animals and fungi, and also done by all other life forms. The organic carbon in this reaction is our food, and the oxygen comes from the air. The products of respiration, CO_2 and water vapor, come out when we exhale.

Photosynthesis serves two purposes in the biosphere. First, it creates carbon in the oxidation state at which its chemistry is versatile enough to build the machinery of life. Second, it stores energy from sunlight in the form of reduced carbon. Some biomolecules have a clear use in the living machine, such as **proteins**, which can be structural or can catalyze chemical reactions, in which case they are called **enzymes**. DNA stores patterns for making proteins, the genetic material, and there are **lipid** or fat molecules in cell membranes like soap bubbles, holding stuff in and keeping stuff out. Other types of biomolecules are more obviously for storage of energy, such as glucose in blood sugar or the fat deposits underneath our skin.

Carbon is the structural backbone of life and the way life stores energy.

Nearly all of the organic carbon produced by photosynthesis is respired sooner or later. Dry wood may hold organic carbon for a few centuries, but really long term preservation requires a special environment such as anoxic peat deposits on land or in the sediments of the deep sea.

Life on Earth has built up a sizable pool of reduced carbon in ocean sediments and **sedimentary rocks** (former ocean sediments much of it currently on land). When oxygen-producing photosynthesis first began on Earth, the oxygen that was released reacted with other reduced species such as dissolved iron in seawater. Sulfur was oxidized to the chemical form of sulfate, SO_4^{2-} , a component of sea salt in today's ocean.

After the biosphere "rusted" as much as it was going to, oxygen started to build up in the atmosphere. As a result of these other chemical reactions, there is enough organic carbon buried in sediments to use up the O_2 in the atmosphere about 10 times. Visionary Earth scientist James Lovelock imagines the biota "charging up" the biosphere, like a giant battery, by building up a surplus of reactive oxygen and reduced carbon. Fossil fuels are a way of extracting some of the energy in Lovelock's battery.

Oxygen in the air and organic carbon in the ground are the poles of the battery of the planetary biosphere.




Figure 8-1 Carbon reservoirs on Earth and carbon fluxes between them.

The CO₂ in the atmosphere is only the tiniest fraction of the carbon on Earth (Figure 8-1). It contains about 700 **gigatons** (Gton) of carbon. A Gton is a billion (10⁹) metric tons, equivalent to 10¹⁵ g. The amount of CO₂, relative to the other gases, is about 390 parts per million (ppm), or 0.039%. For every CO₂ molecule in the air, there are roughly 2,500 molecules of other gases, mainly nitrogen and oxygen. If we were able to gather up all of the CO₂ in the atmosphere and bring it down to the ground, to build a column of pure CO₂ gas at sea-level pressure, that column would be about 3 m high, reaching perhaps some second-story windows. If we were to precipitate the CO₂ into an icefall of dry ice (solid CO₂), spread uniformly over the Earth's surface, the layer would be about 7 mm thick. The atmosphere is a very small, very thin carbon reservoir.

There is also carbon stored at the land surface, in the oceans, in sedimentary rocks, and in the deep interior of the Earth. These other carbon reservoirs all “breathe” CO₂, causing atmospheric CO₂ to vary, naturally, on all sorts of time scales, from yearly cycles to changes over millions of years. The reservoirs interact with each other through the atmosphere, which is like the Grand Central Station of the carbon cycle.

The Land Breathes

Two forms of carbon are associated with the landscape that we live in. The actual living carbon, trees and camels and all the rest of it, is called the **terrestrial biosphere**. About 500 Gtons of C are in the terrestrial biosphere, so it is comparable in size to the atmosphere.

The other land carbon reservoir is the **soil carbon pool**. You can see in Figure 8-1 that it holds about twice as much carbon as the terrestrial biosphere. The organic carbon in soils is mostly dead rather than living: decomposing leaves and other plant material. Organic carbon concentrations in soils are very patchy, depending on the climate, the forestry, and the history of the land, such as fires and agriculture. Deserts do not have much carbon in their soils, whereas grasslands tend to have quite a bit. There is more soil carbon in colder climates than in warmer climates because organic carbon decomposes more quickly when it is warm.




Figure 8-2 CO₂ concentrations in the atmosphere over the past fifty years.

The land breathes on an annual cycle.

The terrestrial biosphere inhales CO₂ from the atmosphere during the growing season and exhales during winter. You can see a seasonal cycle in atmospheric CO₂, especially in the top panel in Figure 8-2. The CO₂ annual cycle in the southern hemisphere is smaller than that in the north because there is less land in the south, and it is out of phase with the north because the seasons are reversed. One interesting observation, foreshadowing Chapter 10, is that the annual cycles in atmospheric CO₂ have been getting larger as atmospheric CO₂ has risen. It is not obvious by eyeballing (Figure 8-2), but the trend emerges when the data are treated statistically. The terrestrial biosphere has been breathing more deeply, as if it is getting larger, with increasing atmospheric CO₂.

The Ocean Breathes

The **ocean carbon reservoir** is larger than either the land surface or the atmosphere. The ocean's carbon is not only dead, it is oxidized: energetically dead as well as biologically dead. The carbon is in the forms CO₂, H₂CO₃, HCO₃⁻, and CO₃⁼, all of which we get to know better in

Chapter 10. The sum of all of them is called **dissolved inorganic carbon**. There is about 50 times as much carbon in the ocean as there is in the atmosphere.

There is also organic, reduced carbon dissolved in seawater; like the organic carbon in soils, it comes from photosynthesis but is no longer alive. Its molecular structure is disordered, a similar chaotic “goo” to organic carbon in soils. And of course there are fish and dolphins and plankton, but the total amount of living carbon in the ocean is small, only about 1 Gton.

Carbon is released from the water to the air in some parts of the world’s oceans, and it dissolves into the water in other places. The total rates of carbon exchange between the atmosphere and oceans are large, comparable to the breathing rates of the terrestrial biosphere and much greater than the rate of anthropogenic carbon release (Chapter 9). The oceans would seem to have the clout to affect the CO_2 in the atmosphere as strongly as the terrestrial biosphere does, and yet the year-to-year breathing variations in the ocean are much less obvious in variations in atmospheric CO_2 . The ocean breathes more steadily than the land does, with less of a seasonal cycle.

Computer models predict that it takes hundreds of years for CO_2 to reach equilibrium between the atmosphere and the ocean. If CO_2 is released to the atmosphere, it will take this long for the ocean to inhale its fill. This is about how long it takes for all the waters to come to the ocean surface somewhere to exchange carbon with the atmosphere.

The ocean affects atmospheric CO_2 on time scales of centuries. The glacial-interglacial cycles were amplified somehow by the ocean carbon cycle.

The clearest example of the power of the ocean to affect CO_2 in the atmosphere is the **glacial-interglacial cycles**, rhythmic and large fluctuations in the amount of ice on Earth in which the climate flips back and forth between glacial states and interglacial states (Figure 8-3). The Earth is currently in an interglacial state called the **Holocene**, and it has been for about




Figure 8-3 CO_2 and methane concentrations in the atmosphere over the past 400,000 years, along with temperature in Antarctica.

the last 10,000 years. This is the stage upon which human civilization has developed, although biologically modern humans evolved considerably earlier.

Before that was a time when it was about as cold as the glacial cycles ever get, called the **Last Glacial Maximum**. The surface of the Earth was about 5–6°C colder than today, and North America and Northern Europe were covered with massive ice domes like what currently exists in Greenland. The ice sheet in American is called the **Laurentide ice sheet**, and the one in Europe is called the **Fennoscandian ice sheet**.

The beat of the ice age rhythm apparently originates from variations in the Earth's orbit around the sun. The orbit varies through three main cycles, each ringing on its own characteristic frequency (Figure 8-4). The first cycle is called the **precession** of the seasons or sometimes precession of the equinoxes (Figure 8-5). The axis of rotation of the Earth spins around like a wobbling top, completing the entire circle in 20,000 years. Most of the solar heat influx variability at high latitudes, and nearly all of the variability in the tropics, comes from precession.

Orbital variations drive climate by changing the distribution of sunlight at the Earth's surface.

Precession has an effect on sunlight intensity because the orbit of the Earth is not circular but elliptical. In its present-day orbital configuration, the Earth is closest to the sun during winter




Figure 8-4 History of orbital forcing of Earth's climate since 400,000 years ago.




Figure 8-5 The precession orbital cycle.

in the northern hemisphere (Figure 8-5). The seasonal cycle in the north is weakened by this orientation because Earth is close to sun when the northern hemisphere is tilted away from the sun. The seasonal cycle in the southern hemisphere is stronger in this configuration because the Earth is close to the sun and tilted toward the sun at the same time. Remember, it is the tilt of the Earth that causes the Earth's seasons (Chapter 6). The precession cycle modifies the seasons somewhat.

Another cycle involves the **obliquity** of the angle of the pole of rotation relative to the plane of Earth's orbit (Figure 8-6). The Earth rotates, making day and night, on a rotation axis of the North and South poles. This rotational axis is not perpendicular to the plane of Earth's orbit around the sun but is tilted somewhat. The angle of tilt is currently 23.5 degrees, but it varies between 22 and about 25.5 degrees on a cycle time of about 41,000 years (Figure 8-4). The impact of obliquity on solar heating is strongest in high latitudes.

The third orbital cycle involves how elliptical the orbit of the Earth is, also called its **eccentricity**. At present, the orbit of the Earth is nearly circular. The eccentricity of the orbit has cycles of 100,000 years and 400,000 years (Figure 8-4). The strongest climate impact of eccentricity is to determine the strength of the precessional forcing. If the Earth's orbit were circular, it would make no difference where the Earth was in its precession cycle because the Earth is equally far from the sun at all parts of the orbit, and the 20,000-year waves in the precession cycle vanish (shaded areas in Figure 8-4).




Figure 8-6 The obliquity of Earth's orbit.

When you average over the entire surface of the Earth and over the entire year, the orbital cycles only have a tiny effect on the total amount of heat the Earth gets from the sun. The orbital cycles affect climate by redistributing the energy from one place to another and from one season to another.

It turns out that the climate of the Earth is especially sensitive to the solar heat flux at about 65 degrees latitude in the northern hemisphere summer. It makes sense that the sweet spot for driving the ice ages is in the northern hemisphere because this is where the ice sheets grow and melt the most through the glacial cycles. The ice sheet in Antarctica grows and melts back somewhat through the climate cycles, but it does not disappear entirely during warm times the way the Laurentide and Fennoscandian ice sheets have. It is summer sunlight that matters most to driving climate rather than winter as you might have expected because it is always cold enough to snow in winter at these latitudes, but the question is whether the summer is warm enough to melt the snow or keep it around to add to the snow from the next winter. Variations in sunlight intensity in June, 65 degrees north latitude, calculated from models of the Earth's orbit, correlate well with the history of the amount of ice on Earth, inferred from ice core and deep sea sediment records.

Summer solstice, 65 degrees north latitude, is the sweet spot for sunlight intensity to drive Earth's climate through the glacial cycles.

Glacial cycles in atmospheric CO₂ were discovered by measuring the CO₂ concentrations in bubbles of ancient atmosphere trapped in the ice sheet of Antarctica. The trapped air bubbles from glacial times have a lower proportion of CO₂ in them relative to the other gases, about 180 to 200 ppm, rising to 260 to 280 ppm during interglacial intervals before the industrial era. The atmospheric CO₂ changes are responsible for about half of the climate changes through the glacial cycle. The changes in the albedo of the Earth, from the large ice sheets and sea ice, are responsible for the other half. The climate during glacial time is an important test of the climate models used to forecast global warming (Chapter 11).

No one is sure exactly why CO₂ in the atmosphere changes up and down in concert with the ice sheets, but the ocean is thought to be the major player. The land carbon reservoirs were, if anything, smaller during the glacial time, which would by itself have left atmospheric CO₂ higher during glacial time than interglacial, rather than lower as the ice core records show. The carbon cycle to and from the solid Earth, discussed in the next section, is too slow to explain the CO₂ changes, some of which are as fast as 1,000 years, which is roughly the time scale for CO₂ to equilibrate between the atmosphere and the ocean. The ocean is thought to be the only pool large enough and potentially reactive enough to explain the CO₂ cycles. This book will not answer the question of how exactly the ocean carbon cycle positive climate feedback worked, but it will discuss ocean carbon chemistry a bit further in Chapter 10.

The Rocks Breathe

The **sedimentary rock carbon pool** is larger still than the ocean, land, or atmospheric pools. Carbon in the solid Earth exists in the form of **limestones**, CaCO₃, and to a lesser extent organic carbon. Most of the organic carbon in sedimentary rocks is in a form called **kerogen**. Kerogen is useless as a fossil fuel because it is dilute, usually less than 1% by weight of sedimentary rocks, and because it is in a solid form, making it difficult to extract. Fossil fuels are discussed in Chapter 9. These carbon reservoirs together contain about 500 times as much carbon as the atmosphere and the landscape combined.

The solid Earth is the largest but slowest breathing of the carbon reservoirs.

Carbon is cycled between the atmosphere and the solid Earth by means of a chemical reaction called the **Urey reaction**, which is



within a planetary geochemical cycle shown in Figure 8-7.

CaSiO_3 on the left-hand side is a simplified chemical formula for a rock typical of the Earth's bedrock, called a **silicate rock**. These usually solidify from melted rock, which is called **lava** if it is found at Earth's surface and **magma** if it is subsurface. Real silicate rocks have other elements in them and a wide range of chemical formulas, but the simple formula CaSiO_3 works for conveying the essential idea of how the rocks interact with the carbon cycle. CaCO_3 and SiO_2 , the solid phases on the right-hand side of the reaction, are typical sedimentary rocks, which form at cool temperatures from elements that were dissolved in water.

The Urey reaction running from left to right, producing sedimentary rocks from silicate rocks, is called **chemical weathering**. A silicate rock weathers by dissolving into river water, ultimately to be delivered to the ocean. Real silicate rocks often contain other elements that do not dissolve in water very well, like aluminum, so these elements form other solids, called **clay minerals**, stuff that you find in soils. After the dissolved calcium and silica reach the ocean, organisms like corals and shell-forming plankton construct solid CaCO_3 and SiO_2 shells from them.

At the cool surface of the Earth, oxidized carbon wants to be in CaCO_3 .




Figure 8-7 Components of the silicate weathering thermostat.

Weathering is the preferred direction for this chemical reaction under the relatively cold, wet conditions of the surface of the Earth. If weathering were allowed to continue to its equilibrium, it would pull nearly all the CO₂ out of the atmosphere.

At high temperatures in the Earth's interior, the Urey reaction can run in the other direction, making silicate rocks from sedimentary rocks. This chemical reaction is called **metamorphic decarbonation**. This direction is generally favored at high temperature, preferring the volatile CO₂ phase the way that water prefers the volatile vapor phase at higher temperature. The CO₂ released by metamorphic reactions may find its way back out of the solid Earth by what is called **degassing**, in volcanic gases or in hot water springs (deep sea vents) at the bottom of the ocean. Some of the carbon degassing from the Earth may be **juvenile** carbon, which has spent the last 4.5 billion years of Earth history bound up in the deep Earth, only to emerge now, but most of it is recycled carbon from previously subducted sedimentary CaCO₃.

In the hot interior of the Earth, oxidized carbon wants to be free, as CO₂.

The fluxes of CO₂ by weathering and degassing are small compared with the other fluxes in Figure 8-2, but if they were out of balance, if you were stop all degassing, for example, weathering would use up all the CO₂ in the atmosphere in a few hundred thousand years. The Earth is much older than this, so if you wait long enough the weathering and degassing fluxes of CO₂ have to balance. Eventually what goes up must come down.

The way the Earth manages to balance the degassing and weathering fluxes is by finding the right CO₂ concentration, a climate-stabilizing mechanism called the **silicate weathering CO₂ thermostat**. The rate of weathering depends on the availability of freshwater rainfall and runoff that rocks can dissolve into. The rate of freshwater runoff depends, in turn, upon the climate of the Earth. In a cold climate there is less overall rainfall, and therefore a slower rate of CO₂ uptake by weathering reactions. If the climate is too cold, an imbalance in the fluxes would allow CO₂ to accumulate in the atmosphere, warming the climate, increasing the weathering rate, and pushing the degassing and weathering fluxes back toward balance. If the climate is too warm, CO₂ is consumed by weathering faster than it is degassed from the Earth, cooling the planet and relaxing towards carbon flux balance.

The CO₂ thermostat regulates atmospheric CO₂ and climate on geologic time scales of hundreds of thousands of years.

This is another opportunity to apply the kitchen sink analogy first used in Chapter 3. Water inflow through the faucet represents CO₂ degassing from the Earth, and flow down the drain, representing weathering, depends on the amount of water in the sink, the stand-in for atmospheric CO₂. The CO₂ in the air drifts toward the value at which the budget balances, just like the water level in the sink does.

The existence of a natural thermostat mechanism that controls the CO₂ concentration of the atmosphere sounds like good news for humanity with our tendency to add CO₂ to the air, but the catch is that it takes hundreds of thousands of years for the thermostat to stabilize CO₂ and climate. On time scales shorter than that, it is perfectly possible for things to be out of whack, such as through the glacial cycles or as a result of fossil fuel CO₂ release. These would be like dumping buckets of water into the sink, or removing water, perturbing the water level for a while until it recovers.

Just like the kitchen sink, the equilibrium of the CO₂ thermostat can be tweaked by messing with the faucet or the drain. The faucet represents CO₂ degassing from the solid Earth, which could be stronger or weaker depending on the geometry of the continents and the volcanoes.

Today a large fraction of the global rate of CO_2 degassing is coming from a particular volcano called Mt. Etna in Sicily, which makes it seem as though CO_2 degassing from the earth might be pretty episodic, variable in time.

Today there is more CaCO_3 accumulating on the sea floor than is subducting into the Earth, because more CaCO_3 accumulates in the Atlantic but more sediment subducts in the Pacific and Indian oceans. This means that at some time in the past, the CaCO_3 subduction rate must have been much higher than it is today, because on long-term average the CaCO_3 fluxes have to balance. Subducting more CaCO_3 would release more CO_2 into the atmosphere, driving the thermostat to make the Earth warmer.

It is possible to change the set-point of the thermostat, creating a hot-house world like that of the dinosaurs or an icy world like today.

The drain can also drive water-level changes in the kitchen sink analogy. The drain represents weathering reactions that take up CO_2 . A mountainous terrain weathers more quickly than a flat plain covered in thick soil, because the soil isolates the silicate bedrock from the rainwater that weathering requires. The Himalayas are the most prominent mountainous feature on Earth, uplifted by the collision of India and Asia a few tens of millions of years ago. Paleoclimatologists have suggested that the uplift of the Himalayas might be responsible for our current era of relatively cool climate, as opposed to the hothouse conditions before then, such as the **Cretaceous period** when dinosaurs ruled the Earth, and **Early Eocene optimum** which followed (Figure 8-8), when there may have been 10 times more CO_2 in the air than there is today.

Plants could affect Earth's climate by pumping CO_2 into the gas spaces in the soil. Soil gas CO_2 concentrations can be 10 times higher than in the atmosphere because of plants. Extra CO_2 tends to supercharge the rate of chemical weathering because CO_2 is an ingredient in the weathering reaction, and also because CO_2 is an acid (see Chapter 10), which drives rocks to




Figure 8-8 History of the temperature of the deep ocean, which tells us about changes in global climate since 65 million years ago.

dissolve. Plants invaded the land surface in the Silurian, about 400 million years ago. Carbon cycle models predict that this change in land surface geochemistry might have had an impact on the long-term climate of the Earth.

The terrestrial planets **Venus**, **Earth**, and **Mars** provide a perfect “Goldilocks” parable to end this chapter. The sin of Venus (the planet) was to be too close to the sun. The water that must have originally been packaged with Venus evaporated in a runaway greenhouse effect and its hydrogen was lost to space forever. With no water, silicate weathering reactions come to a halt, allowing Venus’s allotment of carbon to build up in the atmosphere. The thermostat is broken. As a result, Venus’s carbon wound up as CO_2 in the atmosphere.

The thermostats of Venus and Mars are broken.

Mars is small, about half the diameter of the Earth, so that its interior cooled faster than Earth’s. Mars boasts the largest volcano in the solar system, Olympus Mons, but today all of Mars is geologically dead. The partial pressure of CO_2 ($p\text{CO}_2$) on the surface of Mars is actually higher than it is on Earth, but there is not enough CO_2 in the air to bring the planet to the melting point of ice, which would allow chemical weathering reactions to occur. If there were ongoing CO_2 degassing from the planet, and if there were water, presumably eventually the CO_2 concentration would build up to that level. But there is no degassing so the CO_2 thermostat on Mars is broken, but in a different way than it is broken on Venus.

TAKE-HOME POINTS

- The most stable form on carbon on Earth is oxidized, like CO_2 or CaCO_3 . Photosynthesis stores energy from the sun by producing organic carbon, which also serves as scaffolding for the machinery of life.
- There is less carbon in the atmosphere than there is in other carbon reservoirs on Earth, such as the terrestrial biosphere and the oceans. These other reservoirs tug on atmospheric CO_2 , seasonally for the land, and on glacial interglacial 100,000-year time scales from the ocean.
- The weathering of igneous rocks on land controls the $p\text{CO}_2$ of the atmosphere on million-year time scales. This silicate thermostat stabilizes climate. The thermostat is broken on Venus because there is no water left, and on Mars because there is no active volcanism left.

STUDY QUESTIONS

1. Describe the special role of the atmosphere in the global carbon cycle.
2. Describe the multiple roles of the element carbon at or near the surface of the Earth.
3. Write examples of organic and inorganic carbon molecules, and figure out their oxidation states.
4. Is the CO_2 thermostat a positive or a negative feedback? Relate the components of the thermostat to the different parts of the kitchen sink analogy.
5. Does the carbon cycle provide a positive or a negative feedback through the glacial-interglacial cycles?
6. Describe how the CO_2 thermostat mechanism broke on our sibling planets, Venus and Mars.

FURTHER READING

Alley, R. B. *The Two-Mile Time Machine: Ice Cores, Abrupt Climate Change, and Our Future*. Princeton University Press, Princeton, NJ: Princeton University Press, 2002.

Exercises

Answer these questions using an online geologic carbon cycle model located at <http://understandingtheforecast.org/Projects/geocarb.html>. The model predicts the influence of geologic processes on the CO₂ concentration of the atmosphere.

The model does an initial spin-up phase to make sure that the initial condition is an equilibrium where all the various budgets balance. You will see only the last few years of this phase in the model output. Then there is a transition to a new set of conditions, say an increase in CO₂ degassing flux from volcanoes. We can watch the transition from one equilibrium state to another.

At the moment of the transition from spin-up phase to transient phase, you can inject a slug of new CO₂ into the system, say by fossil carbon combustion and release to the atmosphere. For the problems below, we'll set the fossil fuel carbon spike to 0.

1. **Weathering as a function of CO₂.** In steady state, the rate of weathering must balance the rate of CO₂ degassing from the Earth, from volcanoes and deep sea vents. Run a simulation with a higher CO₂ degassing rate at the transition time.
 - a) Does an increase in CO₂ degassing drive atmospheric CO₂ up or down? How long does this take?
 - b) How can you see that the model balances weathering against CO₂ degassing?
 - c) Repeat this run with a range of degassing rates, and make a table of the equilibrium CO₂ concentration as a function of the CO₂ degassing rate. The CO₂ degassing rate is supposed to balance the CO₂ consumption rate by silicate weathering.
 - d) Make a plot of weathering as a function of atmospheric pCO₂ using your model runs.
2. **Effect of solar intensity on steady state CO₂ concentration.** The rate of weathering is a function of CO₂ and sunlight, a positive function of both variables. By this I mean that an increase in CO₂ will drive an increase in weathering, as will an increase in sunlight. The sun used to be less intense than it is now. Turn back the clock 500 million years to when the sun was cooler than today. What do you get for the steady-state CO₂, and how does this compare with what you get for today's solar intensity? Explain why.
3. **Plants.** Plants pump CO₂ down into the soil gas, possibly accelerating weathering. They also stabilize soils, perhaps decreasing weathering. Run a simulation with a transition from no plants to a world with plants, with no carbon spike on the transition, to see which way the CO₂ effect of plants will go.
 - a) What happens to the rate of chemical weathering immediately after the plants are introduced? Does it go up or down? What does it do after some time has passed?
 - b) What happens to atmospheric CO₂ and why?
 - c) When the CO₂ concentration changes, where does the carbon go?

Where fossil fuels come from and how they are used

Energy Sources

The flow of energy drives the biosphere of the Earth, and it drives human civilization. Energy flows can be described in units of **Watts**, equal to **Joules** per second, where Joules is a unit of energy like calories. A hairdryer might use 1000 W of power, or 1 kiloWatt (kW). For global energy fluxes, it is convenient to speak in terms of **teraWatts**, equal to 10^{12} Watts, abbreviated as TW.

The sun bathes the Earth in energy at a rate of about 173,000 TW. Photosynthesis captures about 100 TW. Mankind is consuming about 14 TW of energy per year. Some of our energy use comes from the sun; wind, hydroelectric, solar, and biomass energy are all driven ultimately by sunlight. Of these, the largest source of energy is hydroelectric power, which makes up about 2% of the global energy supply (Figure 9-1). Renewable energy, defined as geothermal, solar, wind, wood, and waste electric power, supplies about 0.3% of our power.

All of our energy ultimately comes from stars, mostly our sun.

Humans are also mining one-use-only energy sources from the Earth, in fossil fuels and radioactive elements. Nuclear energy is mostly driven by decay of uranium, which was formed into its unstable high-energy nuclear state in supernova explosions of stars. I guess you could say that nuclear energy is fossil starlight. It accounts for another 2% of global energy production.

Fossil fuels derive their energy from ancient sunlight, harvested by photosynthesis that took place millions of years ago. Over Earth's history, photosynthesis has preserved a sizable amount of reduced carbon, and the energy in this carbon can be harvested by reacting it back with O_2 to regenerate CO_2 and H_2O . In other words, the carbon can be burned, releasing its energy as heat; or it can be respired, as when you digest your lunch, capturing the energy to drive biological processes.

Most of the reduced carbon on Earth is so diluted in rocks and sediments that for practical purposes it is useless as a fuel. It takes very special circumstances to preserve the carbon in a pure enough form to be a convenient energy source. Fossil fuel carbon comes in all three phases: solids, liquids, and gases.

Only a small fraction of the buried organic carbon is in a convenient form for fuel.

The largest traditional fossil fuel reservoir is **coal**. The carbon in coal comes from land plants, that were deposited in swamps. The organic carbon is preserved in swamps because the standing water protects it from degradation by oxygen in the atmosphere.




Figure 9-1 The sources of energy used globally and in various countries. From the U.S. Energy Information Administration, data from 2006 and 2007.

Coal from freshwater swamps contains less sulfur than coal from saltwater swamps, because of the sulfate, SO_4^{2-} , dissolved in seawater that winds up as sulfur in the coal. Upon combustion, the sulfur forms droplets in the air called *aerosols* that scatter sunlight, then it deposits as acid rain. High-sulfur coal also has mercury in it, making coal power a leading source of mercury release to the environment.

Coal is the most abundant fossil fuel, and the future of Earth's climate depends mostly on what happens to that coal.

Over time, the carbon accumulates into thick layers of what is called **peat**. If the peat is buried and cooked for millions of years, it turns into coal. There are different grades of coal depending on how well cooked it is. The highest grade of coal has been cooked the longest, and it is composed almost entirely of carbon, the original hydrogen and oxygen driven off as water, CO_2 , oil, and gas.

The British Petroleum Company (BP), in their annual report on global energy (see Further Reading) estimates that there are 1,000 gigatons (Gtons) of carbon in coal reserves. Another source, an academic article by Rogner (also listed in Further Reading), estimates that the total amount of coal ultimately available might be 10 times that. David Rutledge at CalTech argues that there is much less coal available than this. However, it is commonly assumed that there is enough coal to supply our needs for several centuries, if we choose to use it.

In the United States, coal is mostly used for generating electricity (Figure 9-2). Coal itself is cheap and abundant in many parts of the world, and if a coal-fired power plant is already built, it can produce electricity more cheaply than pretty much any other source.

The energy stored in **oil** was originally harvested from sunlight by photosynthesis that took place in the ocean. Plants in water are generically called **algae**. Microscopic algae suspended in the ocean are called **phytoplankton**. Phytoplankton in the ocean produces about as much organic carbon by photosynthesis every year as the terrestrial biosphere on land does. A tiny fraction of the dead phytoplankton ends up in sediments; the rest gets eaten by something, animal or bacteria, in the water column or in the top layer of sediments.

The sediments covering most of the sea floor do not contain enough organic carbon to ever make oil. Typical deep-ocean sediments have less than 1% organic carbon by weight, but typical **oil source rocks** contain something like 3% to 10% organic carbon. Most oil comes from very




Figure 9-2 How different forms of fossil energy are used in the United States. From the U.S. Energy Information Administration.

specific oil-source rock layers just a few feet thick. These source rocks formed near continents and under water that has no oxygen dissolved in it, which is called **anaerobic** conditions.

If organic-rich sediments are buried to a depth in the Earth between 2 and 5 km, they will be heated to temperatures of 50° to 150°C. This has the effect of converting some of the dead plankton into oil, consisting of chains of carbons with hydrogens attached (Figure 9-3). If the temperature goes higher, the carbon winds up as methane in natural gas.

Oil is probably the most convenient but the least abundant of the fossil fuels, so it is the most expensive.

If the mud is too impermeable, any oil or gas produced could just stay there in small droplets or bubbles, out of drilling reach. Alternatively, it could flow through the pores of




Figure 9-3 The structure of normal octane, a hydrocarbon chain such as found in petroleum.




Figure 9-4 The amount of proven petroleum reserves, by region, from British Petroleum.

sandstone or through faults and cracks in the sediment column, perhaps to collect in some upside-down pool in an upside-down bowl of impermeable rocks. Only a tiny fraction of the oil and gas ever produced is in a harvestable form today, that is to say, mobile enough to have collected together but not mobile enough to get squeezed all the way out of the Earth.

Oil is the most expensive of the fossil fuels because it is uniquely suited to fueling transportation (Figure 9-2). As a combustible liquid, it easily carries more energy per weight than any batteries that have been developed so far. Coal has not been used for transportation since the days of the steam locomotive, and gas must be carried in a pressurized container, as opposed to petroleum, which sits happily in an unpressurized fuel tank.

Perhaps because such special conditions are required to produce harvestable oil, the distribution of oil deposits on Earth is very “spotty” (Figure 9-4). Most of the oil is in the Middle East, although there is a smattering scattered throughout the rest of the world. We do not know precisely how much oil there is available to harvest. The BP report tallies the proven reserves of oil, stuff that has been discovered and documented already. They say 1150 thousand million barrels of oil, which translates to 135 Gton C. The amount of oil ultimately extractable will almost certainly be higher than this. Technological advances in oil extraction can double the amount of oil you can get from an oil field. More oil fields will certainly be discovered, although there are limits to how much more oil we can expect to find. Half of the world’s oil reserves today are in the 100 biggest oil fields, and most of the largest ones were discovered decades ago.

There is a class of carbon deposits for which the geological conditions were never quite right to produce oil, but from which oil can be extracted if the rocks are mined and cooked, or cooked in place by injecting steam into the well. These nontraditional oil sources are called **oil shales** and **tar sands**. The Canadians are mining and processing tar sands now. If future discoveries, technological progress, and nontraditional sources are counted together, the estimates disagree wildly, but some of them begin to reach 500 Gton C for how much “oil” there might ultimately be.

How long will the oil last? One way to calculate a number with units of time would be

$$\text{Lifetime}[\text{yr}] = \text{reserves} [\text{kg}] \cdot \frac{1}{\text{production rate}} \left[\frac{\text{yr}}{\text{kg}} \right]$$

British Petroleum calls this the **reserves to production** or **R/P** ratio. Notice that the units on both sides of the equation balance, years on both sides. The R/P ratio tells how long it would take to use up a resource if we continued using it at the rate we are using it now. The funny




Figure 9-5 The value of the production/reserves ratio over the last 25 years. From British Petroleum.

thing about this ratio is that its value has not changed in 15 years (Figure 9-5). According to BP figures, we have 40 years of oil left, and it looks like we always had 40 years left, and we could have 40 years left forever. The clock seems to have stopped. Actually, what is happening is that oil is being discovered, and extraction efficiency is increasing, about as quickly as oil is being used. Also, the amount of oil that is classified as reserves depends on the price of oil because it only counts if it is currently profitable. The rate of consumption is growing with time, but we have been adding new oil to the reserves quickly enough to more or less keep pace with that growing demand.

There is enough oil to keep pumping for decades, but the peak rate of oil extraction could be happening right now.

Another way to estimate the lifetime of the age of oil is called **peak oil**, developed by geologist M. King Hubbert (see the Deffeyes books listed in Further Reading). Hubbert pointed out that the rate of extraction of a limited natural resource such as oil tends to follow a bell-shaped curve (Figure 9-6). The industry has a spin-up time at the beginning, gets going faster and faster, and then slows down as the resource starts becoming depleted. The crucial point is that the maximum rate of extraction comes when the resource is about half used up. Our oil is about half used up, so peak oil theory predicts that the peak is soon upon us.

This is an example of an **empirical** observation; there is no theoretical underpinning to the curve, although it sure does fit nice. But one could imagine an “extract-the-juice-from-a-popsicle-on-a-hot-day curve.” The popsicle consumption rate starts off slowly at the beginning because the popsicle is too cold to eat; then, at the very end, you have to swallow the last half of the popsicle in one gulp to keep it from hitting the sidewalk (Figure 9-7). Oil analysts like Daniel Yergin (see Supplemental Reading section) argue that political and economic factors could modify the oil extraction trajectory to look more like my popsicle curve or some other.

This could be, but a bell curve seems to have worked in the past for individual oil fields. Hubbert used his idea, in 1956, to forecast the peak extraction of oil in the United States. He




Figure 9-6 Oil production in the United States follows a bell curve. M. King Hubbert predicted in 1956 that the peak would come at around 1970, and he turned out to be correct.




Figure 9-7 A bell curve and an “eating-a-popsicle-on-a-hot-day” curve.

predicted the peak to be sometime between 1965 and 1972, and in fact it came in 1970. The solid line in Figure 9-6 shows the time up to Hubbert’s forecast, and the dashed line shows a prediction like his for the future (benefiting a bit from hindsight, I confess). Hubbert’s prediction came at a time when oil production in the United States seemed endless. Figure 9-8 shows a Hubbert’s Peak for sperm oil harvesting back in the whaling days.

The forecast for when the peak in world petroleum extraction will come depends a lot on what you call oil and how much is ultimately out there. Figure 9-9 shows world oil production in the past, fit to two **Gaussian** (bell) curves, one for 200 Gton C in oil that will ultimately be extractable and one for 500 Gton C. Humans have already extracted 117 Gton C. If 500 Gton C is ultimately correct, then the peak comes sometime in the 2030s. If 200 Gton C is correct, consistent with BP, the peak should be happening now. And, in fact, most oil producers are pumping oil as fast as they can now, although oil analysts argue that logistics and infrastructure are slowing things down, in Iraq for example, or in the number of refineries that are operating.

The crucial feature of the bell curve is that the maximum rate of extraction, the peak, occurs when half of the resource is still in the ground, as opposed to the peak of the popsicle-eating




Figure 9-8 A Hubbert curve for the rate of sperm oil production. With permission from U. Bardi, Università di Firenze.




Figure 9-9 Global oil production, history so far plus two Hubbert's Peak projections.

curve, which comes at the very end. During the first part of the extraction, the growth in supply rate follows along with an exponential growth in demand, driven by increasing population and gross domestic product. If the production rate stops growing at the peak but demand continues to grow, it will appear to the market as a shortage. The lesson of peak oil theory is that the reserves to production ratio from BP gives a misleading impression, because the shortage would begin long before the oil is down to the last drop.




Figure 9-10 Natural gas emits less CO₂/W of energy than oil or coal.

The other main type of fossil carbon is **natural gas**, most of which is the gas **methane**. Methane carries more energy per carbon than the others because methane is the most chemically reduced form of carbon. You get energy from making CO₂ from the carbon from all fossil fuels, but you also get energy from converting hydrogen to H₂O. Methane, being more reduced, has more hydrogen than coal or oil does. Turning it around in Figure 9-10, a Watt of energy from methane releases less carbon than a Watt from coal or wood.

Methane also burns more cleanly than the other fossil fuels, producing none of the acid rain, mercury, and aerosols from coal or the smog from petroleum combustion. Methane is more difficult to transport than liquid petroleum because it is a gas and must be held under pressure. For this reason, methane associated with oil in remote locations is sometimes just burned, dramatic flares into the sky, rather than collected. Natural gas prices vary widely around the world, and they tend to fluctuate in time more than, say, coal prices do.

Methane is the cleanest fossil fuel, and it carries the most energy per CO₂ produced.

The amount of extractable natural gas on Earth is not known as well as it is for petroleum because gas has not been explored for as intensively. Industry estimates of available gas are similar to the amount of oil, hundreds of Gton C, as proven reserves. Recently developed techniques for fracturing, or “fracking” the rock to release gas have revolutionized, perhaps doubled, the amount of gas within our reach. Like for oil, the ultimately available natural gas reservoir will certainly keep rising.

Under the right conditions, methane can freeze with water into a **methane hydrate** or **clathrate** structure, soccer-ball cages of water ice, each surrounding a single molecule of methane. Most of the methane hydrate is in mid-depth sediments of the ocean, but some can be found on land in permafrost soils. The deep sediment column is sort of a backwater in the carbon cycle, where life and chemical reactions move very slowly. But over millions of years, thousands of Gton C of methane have built up, swamping by a factor of 10 the other reservoirs of swamp gas on Earth. There is enough methane down there to totally clobber the climate

of the Earth, if it were to be released all at once, because methane is such a powerful greenhouse gas.

The methane in the hydrates seems intrinsically precarious in the carbon cycle. A model of hydrates predicts that there would not be much frozen hydrate if the ocean was about 3°C warmer, as warm it could eventually get if the climate of the surface warmed from fossil fuel CO₂. Fortunately, it takes thousands of years for the heat from the surface to reach the hydrates.

Methane hydrates could be a slow positive feedback in the global carbon cycle, but they will probably not burp out all at once.

Melting hydrate could destabilize the sediment column to submarine landslides by producing bubbles that lift the mud grains off of each other. Hydrate ice is like regular water ice in that it would float in water if it were not buried in mud. So it is possible that submarine landslides could release methane, some of which might reach the atmosphere. Fortunately again for us, the amount of methane that could be released in a single landslide is not enough to really change the climate.

Methane hydrates are not traditionally thought of as a fossil fuel because they are difficult to extract; this is because they are frozen and do not flow unless you melt them and because they are dangerous, leading potentially to submarine landslides. But fuel companies are exploring ways of extracting methane from hydrates on land, which may end up supplying some modest fraction of our natural gas consumption.

The bottom line is that coal is the most abundant available fossil fuel, 10 to 50 times as abundant as either traditional oil or natural gas. The climate impacts of releasing various amounts of carbon are discussed in Chapter 13, but the conclusion there will be that coal is the form of fossil fuel with the capacity to take the climate of the Earth well beyond a sort of turning-point temperature change of 2°C. The future of Earth's climate will depend most on what happens to that coal.

Energy Consumption

Different parts of the world get their energy from very different portfolios of sources (Figure 9-1). India and China, the largest of the developing world, use a lot of coal, whereas Brazil uses mostly petroleum and hydroelectric power. Denmark is remarkable for windmills, which fall into the category “renewable” on this plot. France has invested heavily in nuclear energy (not renewable). The United States energy sources are similar to the global average. In part this is because the United States is a big consumer of energy, accounting for about a quarter of energy use globally (Figure 9-11).

Dividing the energy use of a country by the number of people in it gives the energy use per person or **per capita**. China has recently overtaken the United States as the country emitting the most CO₂, but there are many more Chinese people than there are Americans so that the average Chinese person emits one-tenth the amount of carbon released by the average American. The Europeans and Japanese only use half the energy per capita as the United States. Energy use in the U.S. is five times higher than the global average, so that if everyone adopted the American lifestyle it would increase the global rate of carbon emissions by about a factor of five.

Energy use is closely related to economic productivity, which is measured by the total value of everything produced for sale in a country in a year, the gross domestic productivity or




Figure 9-11 Energy consumption (*top panel*), energy consumption per person (*middle*), and energy consumption per dollar gross domestic product (GDP) (*bottom*), globally and from various countries. Data from 2008 and 2009.

GDP, units of dollars/year. Dividing energy use by GDP gives the **energy intensity**, a measure of the energy efficiency of an economy, reflecting waste or efficiency but also the difference between heavy industry such as steel production versus high-tech or service economies. Figure 9-11 shows that the Europeans and the Japanese can make money using less energy than the United States can, while the Chinese and Indians are more energy intensive (less efficient). The patterns of carbon emissions from these countries are similar to the patterns of energy use (Figure 9-12). The United States is responsible for an about a quarter of global CO₂ emissions.

Energy use and carbon emissions are related to population, economic activity, efficiency, and energy sources.

The factors that control CO₂ emission can be combined together to predict carbon emissions in the future if there were no effort to stop: a scenario known as **business as usual**. The




Figure 9-12 Carbon release in CO₂ per year (*top panel*), carbon release per person (*middle panel*), and carbon release per dollar gross domestic product (GDP) (*bottom*), globally and from various countries. Data from 2008 and 2009.

equation is called the **Kaya Identity**, and it is the basis for one of the online models used in the Exercises section. The results from the model are shown in Figure 9-13. The Kaya Identity is

$$\begin{aligned}
 \text{CO}_2 \text{ emission } \left[\frac{\text{kg C}}{\text{year}} \right] &= \text{Population}[\text{persons}] \\
 &\quad \times \text{GDP per capita} \left[\frac{\$}{\text{person}} \right] \\
 &\quad \times \text{energy intensity} \left[\frac{\text{Watts}}{\$} \right] \\
 &\quad \times \text{carbon efficiency} \left[\frac{\text{kg C/year}}{\text{Watts}} \right]
 \end{aligned}$$

Check it out for yourself that the units cancel correctly.




Figure 9-13 Results from the Kaya Identity model (<http://understandingtheforecast.org/Projects/kaya.html>).

Population has been increasing at a rate of about 1.3% per year, but this rate of growth is not expected to continue forever. Rates of population growth tend to decrease with the affluence of societies. Population is forecast to level off at some value, but that value could be 9 billion people or it could be 15 billion people. A typical “business-as-usual” forecast puts the leveling-off point at about 11 billion people.

The second factor, \$GDP/person, has been rising throughout the last century at a rate of about 1.6% per year, rising from \$930 in 1900 U.S. dollars in 1900, to \$4,800 per person in 2000.

The third factor is the energy intensity, how many watts of energy it takes to make \$1 GDP. The energy intensity reflects efficiency but also the difference between heavy and light industry (Figure 9-12, bottom). Energy intensity has been dropping by a rate of about 0.55%/year over the past century, but the last few decades have seen a drop by about 1% per year. Assuming that our energy efficiency will continue to improve at 1% per year, the Kaya identity model on the website predicts that atmospheric $p\text{CO}_2$ will reach 700 ppm by the year 2100 (Figure 9-14).

The final factor, the carbon released per energy yield, is the result of efficiency of the power plants but also on the characteristics of the feedstock energy source. Coal contains intrinsically less energy per carbon than does oil or gas, whereas nuclear and hydroelectric energy do not result in much carbon release, if any.

The Kaya model can be used to predict how much carbon will be released in the future, and what it would take to change those carbon emissions. You can compare your results to the CO_2 emission scenarios from the Intergovernmental Panel on Climate Change, or IPCC, which ranged from 6 to 26 Gton C emission per year in 2100.




Figure 9-14 More results from the Kaya Identity model.

Chapter 10 describes what happens to that CO_2 after it is released, and Chapter 13 discusses what it means, in terms of energy, to avoid as much as possible the climate impacts of fossil energy.

TAKE-HOME POINTS

- Ultimately, the energy available to humankind includes instantaneous solar energy, which is abundant but spread out; stored solar energy in the form of fossil fuels; and stored energy from stellar explosions in the form of radioactive elements.
- Of the fossil fuels, coal is the most abundant. Oil may run out in the coming decades, and the peak rate of oil extraction may be upon us even now.
- We can project energy demand in the future as the product of population, economic growth, and energy efficiency.

STUDY QUESTIONS

1. Contrast the three forms of fossil fuel carbon, in terms of their origins, inventories, and uses.
2. Which are the least and most reduced forms of carbon fuel? Which form carries the most energy per carbon?
3. Compare peak oil theory to last-drop R/P calculations. What are the strengths and weaknesses of each? How does each relate to economic shortage of oil?

FURTHER READING

BritishPetroleum. Energy in Focus. *BP Statistical Review of World Energy*, 2009.

Deffeyes, K. S. *Hubbert's Peak: The Impending World Oil Shortage*. Princeton, NJ: Princeton University Press, 2003.

Rogner, H. H. An assessment of world hydrocarbon resources. *Annual Review of Energy and the Environment*. 22:217–262, 1997.

U.S. Department of Energy, Energy Information Administration. *International Energy Outlook 2003*. Available at: www.eia.doe.gov.

Exercises

1. **Hubbert's Peak.** Point your web browser to <http://understandingtheforecast.org/Projects/hubbert.html>.
 - a) You will see two different data sets to plot against, along with the three parameters (knobs) that control the shape of the Hubbert curve. Nothing fancy here: We are just matching the curve to the data by eye. First start out with the U.S. oil production. The page comes up with some values for the curve that look pretty good to me, but you should try varying the numbers in the box to see how tightly those values are constrained. In particular, there may be combinations of values that could be changed together to fit the data nearly as well but with different values. How much wiggle room is there for U.S. oil production?
 - b) Now switch to global oil production with the pull-down menu. When do you forecast the peak of oil extraction? How does it depend on your assumption of how much oil will eventually be extractable. How much wiggle-room is there for the year of the peak in global oil extraction?
2. **The Kaya Identity.** Point your web browser to <http://understandingtheforecast.org/Projects/kaya.html>.
 - a) Find the plots for gross domestic product per capita, energy intensity, and carbon efficiency, and compare the model hind-cast (*solid line*) with the data (*plus signs*). How well constrained are the growth rates by the past data? Make a table showing the high and low values that

you find plausible for each parameter. Of course, the future might not follow the dictates of the past; this is not a mechanistic prediction but just a blind extrapolation. Using the past as a fallible guide, however, take a guess at what the range of possibilities is for each of the input values. Remember that these assume business as usual, with no decrease in carbon emissions to avoid climate change.

- b) How much carbon is humankind predicted to emit by the end of the century? For each of the input parameters that you estimated the uncertainty of in part a, which end of the uncertainty range gives the highest and lowest carbon emissions? Add columns to your table showing the CO₂ emission rates for the high and low parameter estimates. Combining the uncertainties in all of the parameters, run the model with all the high-CO₂ emission values of the parameters and then the low-emission values, to find the highest and lowest plausible carbon emission for 2100. The emissions are used to drive a carbon cycle model, called ISAM (**I**ntegrated **S**cience **A**ssessment **M**odel), the results of which are shown in the second output plot. According to the ISAM model, what could the range in atmospheric pCO₂ be like then?
- c) What rate of improvement in carbon efficiency would be required to limit atmospheric CO₂ to at most 450 ppm?
- d) Politicians often talk about CO₂ cuts as a fraction cut by the year 2050. The rate of carbon emission as a function of time would decrease with time following an exponential function according to

$$\text{fraction cuts} = \text{emission rate}(\text{time}) / \text{emission rate}(0) = e^{kt}$$

where

$$k = x\% / \text{yr} / 100\%$$

if x is the rate of change of carbon efficiency, like the -0.3% per year that comes up in the model interface by default. Using the carbon efficiency improvement rate from part c, calculate the fraction cut from 2000 to 2050 required to stabilize CO₂ at 450.

- 3. **IPCC CO₂ Emission Scenarios.** Open a new browser window for the ISAM carbon cycle model, at <http://understandingtheforecast.org/Projects/isam.html>. This page shows the results of IPCC carbon emission scenarios made by more sophisticated crystal balls than our simple Kaya identity (and then offers to run the results of the scenarios through a carbon cycle model). On the Kaya page, try to reproduce the year-2100 carbon emissions from scenarios A (business as usual, or BAU), B (BAU with carbon stabilization), C (slow growth), and F (gonzo emissions). What input parameters are required?

The atmosphere ain't what it used to be

Ozone

Before we get to global warming, it's useful to consider two similar issues: ozone and atmospheric methane. **Ozone** is a reactive oxygen molecule composed of three oxygen atoms. Ozone in the stratosphere is produced when O_2 molecules are zapped by energetic ultraviolet (UV) light, called UV-C, and break apart into two very reactive oxygen atoms. Each of these O atoms may find another O_2 molecule and join it, to form O_3 , or ozone. Ozone itself absorbs UV light called UV-B that is less energetic per photon but more abundant than the UV-C required to break up an O_2 molecule. Stratospheric ozone filters UV-B radiation that might otherwise reach the surface, causing skin cancers and sunburn.

The ozone hole is a separate issue from global warming.

Atmospheric chemists predicted that ozone concentrations in the stratosphere might be depleted by chemical reactions accompanying the breakdown of chemicals called chlorofluorocarbons, or Freons, inert chlorine-bearing compounds that are used in refrigerators and air conditioners. The prediction was that stratospheric ozone concentrations would gradually decrease over decades. This has been observed, but what scientists did not anticipate was a surprise called the **ozone hole**. The ozone hole is located in the southern hemisphere over Antarctica. In wintertime in this region, it gets cold enough that an unusual form of cloud forms in the stratosphere, composed of frozen particles of nitric acid (HNO_3). It turns out that the frozen nitric acid clouds convert the chlorine from breakdown of Freons into a very reactive form that does not just deplete ozone by a small amount but rather consumes it entirely within that air mass as soon as the sun rises in the Antarctic spring. The ozone hole was first observed in measurements of ozone concentration made by hand, on the ground. After these measurements, old satellite ozone records were examined, and it turned out that the satellite had been seeing the ozone hole for several years, but a data quality algorithm had been programmed to throw out any data that violated common sense, and out with the bad data went the ozone hole. This is an indication of how much of a surprise it was.

The ozone hole is a terrific example of a **smoking gun**, clear-cut proof of the detrimental effects of chlorofluorocarbons on stratospheric ozone. Before the ozone hole was detected, theory predicted a gradual decrease in stratospheric ozone globally, superimposed on natural variability: a difficult prediction to decisively test. As a result of the ozone hole, and the availability of economically and technologically viable alternatives, many countries of the world ratified the Montreal Protocol in 1987, phasing out further production and release of chlorofluorocarbons to the environment.

A “smoking gun” makes it easier for a society to decide to act.

Ozone is also produced by the reaction of gases from industrial activity, mostly automobile exhaust, in surface urban smog. The ingredients for ozone production in urban air are evaporated organic carbon compounds, like evaporated gasoline or molecules that are emitted naturally from trees and plants, nitrogen oxide compounds that are emitted from automobiles, and sunlight. When ozone concentrations exceed a toxic limit, asthma sufferers begin to feel discomfort and plant leaves get “burned” and scarred.

Ozone in the stratosphere interacts with climate change caused by greenhouse gases such as CO₂. The main role that ozone plays in the energy budget of stratospheric air is that it absorbs ultraviolet light from the sun, heating up the air. Ozone is the reason why there is a stratosphere. If ozone did not heat up the stratosphere, the atmosphere would continue to get colder with altitude, and the atmosphere would be one big troposphere. The effect of decreasing ozone concentration is less heating and, therefore, cooler air in the stratosphere.

Rising CO₂ concentrations also cool the stratosphere. The role that CO₂ plays in the heat balance of stratospheric air is that it acts like a radiator fin, shedding heat energy as infrared radiation. For this reason, an increase in CO₂ concentration in stratospheric air causes the temperature there to go down. Ozone is also a greenhouse gas, but ozone has a stronger influence on the heat balance of the stratosphere by absorbing UV. So the stratospheric is cooling for two reasons: falling ozone concentration and rising CO₂.

Good ozone and bad ozone.

Ozone is confusing to the student of the environment because ozone in the stratosphere is a good thing, but industrial activity is acting to deplete it; ozone in the troposphere is a bad thing, and industrial activity tends to produce it.

Methane

Methane is a greenhouse gas, about 40 times more powerful per molecule than CO₂ at current concentrations (see Chapter 4). Methane is released to the atmosphere from natural sources and anthropogenic sources (Table 10.1). Once released to the atmosphere, the methane slowly “burns”, turning into CO₂ and H₂O in about a decade.

Methane is a short-lived gas in the atmosphere. This means that the methane concentration would respond quickly to decreasing the human methane sources.

The chemical reaction that spells the end for methane involves a very reactive molecule called OH radical. OH is a water molecule missing a hydrogen atom, and in fact it would much

Table 10.1 Natural and Anthropogenic Sources of Methane (Gton C/year).

Natural sources	
Wetlands	0.075–0.15
Termites	0.015
Human sources	
Energy	0.075
Landfills	0.03
Ruminant animals	0.075
Rice agriculture	0.05
Biomass burning	0.03

prefer to be a water molecule, which is to say, water is a more stable, lower-energy molecule than OH is. So the OH radical steals one hydrogen from methane, CH₄, leaving CH₃, another very reactive and unstable molecule that quickly reacts further.

The OH radical is like the flame at the wick of a lit candle. In the absence of sunlight, there is no OH, and the flame is out. Methane and O₂ gas can coexist together for hundreds of thousands of years, for example in bubbles of ancient air trapped in ice sheets, like a candle can sit in oxygen forever and not decompose if there is no flame. But in sunlight, methane and oxygen slowly burn up.

The methane cycle in the atmosphere is another application of the kitchen sink analogy, already used to describe the energy budget of the surface of the Earth in Chapter 3 and the weathering CO₂ thermostat in Chapter 8. The water level in the sink represents the methane concentration in the air. Inflow from the faucet stands for the production of methane from natural or human sources, which is described below. Flow down the drain is the chemical breakdown of methane in the atmosphere. As with the other kitchen sinks in this book, the flow down the drain (reaction with oxygen) depends on the amount of water in the sink (the concentration of methane in the atmosphere).

The kitchen sink analogy — yet again!

The kitchen sink analogy has an equilibrium where the water flow out balances the input, and so does the methane concentration in the air. The budget that has to be balanced for methane is

$$\text{emission} \left[\frac{\text{Gton C}}{\text{year}} \right] = \text{decomposition} \left[\frac{\text{Gton C}}{\text{year}} \right]$$

Assuming that the decomposition rate is proportional to the concentration gives

$$\text{decomposition} \left[\frac{\text{Gton C}}{\text{year}} \right] = \frac{\text{inventory} [\text{Gton C}]}{\text{lifetime} [\text{year}]}$$

Substituting emission for decomposition,

$$\text{emission} \left[\frac{\text{Gton C}}{\text{year}} \right] = \frac{\text{inventory} [\text{Gton C}]}{\text{lifetime} [\text{year}]}$$

which can be rearranged to

$$\text{inventory} [\text{Gton C}] = \text{CH}_4 \text{ emission} \left[\frac{\text{GtonC}}{\text{year}} \right] \times \text{lifetime} [\text{years}]$$

The inventory is in units of Gton C, which can be converted to the more usual methane units of concentration (mixing ratio, see Chapter 4) in our atmosphere as

$$1 \text{ Gton C in our atmosphere} \approx 2 \text{ ppm} = 2000 \text{ ppb}$$

The methane cycle in the atmosphere has an equilibrium state like the sink does, but like the sink, it does not always have to be in that equilibrium state. The water level in the sink could be rising or falling, although it cannot do this forever because it should be rising or falling toward the equilibrium state, and when it gets there, the water level stops changing.

If you start with the sink at equilibrium and dump a bucket full of water into the sink, the water level will spike upward and then relax back toward the equilibrium level, and the amount of time it takes to do that depends on the lifetime of the water in the sink. The same is true for methane; if there were a sudden release of a slug of methane, the concentration would spike upwards and relax back on a time scale of the methane lifetime in the atmosphere, which is about a decade.

The other way to change the water level in the sink is to alter the equilibrium by changing the flow in through the faucet or by changing the drain in some way, say by partially obstructing it. The water level in the sink will drift toward the new equilibrium and then remain there.

For methane, changes in methane sources are thought to be driving most of the rise in atmospheric methane owing to human activity. There were natural sources of methane to the air before humans arrived, but human activity is adding more, resulting in higher concentrations of methane in the air. The lifetime of methane before it is oxidized also changes somewhat as the methane concentration goes up.

The rise in atmospheric methane is mostly driven by increasing emissions rather than a changing atmospheric lifetime.

One of the natural sources of methane to the atmosphere is the degradation of organic carbon in freshwater swamps. Organic carbon first consumes O_2 , as discussed in Chapter 8. In seawater, after O_2 is gone, organic carbon reacts with sulfate ion, SO_4^{2-} , to produce hydrogen sulfide (H_2S). After oxygen and sulfate are depleted, organic carbon degradation can produce methane. This is how the methane is produced that freezes into clathrate deposits below the sea floor.

In fresh water, there is not much SO_4^{2-} , so, methane production begins as soon as oxygen is gone. Methane is found much shallower in fresh water mud than in saltwater mud. If you step in mucky swampy freshwater mud, you may see bubbles of methane rising up around your legs. This is why methane is known as “swamp gas”.

Bubbles of ancient atmosphere preserved in ice cores tell us that the methane concentration has fluctuated with the climate state over the past 650,000 years (the longest ice core yet available), with lower methane concentrations during colder, drier climate stages (Figure 8-3). This is thought to be the result of the changing abundance of swamps.

Rice is grown in artificial swamps called rice paddies. The rice plant does not, strictly speaking, require the swampy conditions, but the water chokes off some of the competition from other plants, so rice paddies are typically flooded. Other human-caused sources of methane include production in the guts of ruminant animals and leakage by the fossil fuel industry.

The lifetime of methane in the atmosphere seems to have been about constant through the past decades. This is in spite of the fact that there have been significant increases in the abundance of some chemicals in the atmosphere that produce OH radical, like ozone and some nitrogen radical compounds in urban smog. The concentrations of molecules that consume OH radical, methane and carbon monoxide (CO), have also increased, which change by itself ought to have decreased the amount of OH in the air. It seems as though these two types of chemical changes in the air have just about balanced out, so that the methane lifetime has been holding steady.

Overall, as a result of the human impact on the atmosphere, the methane concentration has doubled over its pre-anthropogenic concentration (Figure 10-1) and is responsible for a quarter of anthropogenic greenhouse heat trapping (Figure 10-2).

Carbon Dioxide

CO_2 in the atmosphere differs from methane in that the sinks for CO_2 are much more complicated. CO_2 dissolves in the oceans and ultimately goes back to the solid Earth by means of the weathering CO_2 thermostat discussed in Chapter 8. Methane is an unstable, transient gas in the atmosphere, whereas CO_2 just builds up in the atmosphere-ocean-land surface carbon cycle, because the weathering reactions that provide the ultimate sink are so slow that on human time scales they don't draw down the CO_2 concentration much at all.

Methane is transient, but CO_2 accumulates.




Figure 10-1 History of CO₂ and CH₄ concentrations in the atmosphere from ice cores (symbols) and atmospheric measurements (solid lines after 1950). Replotted from the Intergovernmental Panel on Climate Change, 2001.




Figure 10-2 The impact of various climate drivers on the energy budget of the Earth, in watts per square meter, relative to the year 1750. Replotted from the Intergovernmental Panel on Climate Change, 2007.

There are two main anthropogenic sources of CO₂ to the atmosphere. One is deforestation. A dense forest holds more carbon per area than does a plowed agricultural field. The world's forests started to feel the ax thousands of years ago with the development of agriculture and the growing of the human population. Most of the temperate latitudes have been cut long since, but deforestation in the tropics is still going on today, accounting for about 1.5 Gton C/year release of CO₂ to the atmosphere.

Atmospheric CO₂ started to rise, according to measurements of the chemistry of bubbles in ice cores, about the year 1750, after millennia of holding steady at a natural value of about 280 ppm (Figure 10-1). For the first century or so after that, most of the concentration increase was probably caused by deforestation in the New World (the “**pioneer effect**”). Today ongoing deforestation is mostly to be found in the tropics.

Cut down a tree, and its carbon will eventually go to the atmosphere.

The other main anthropogenic CO₂ source is, of course, the combustion of fossil fuels discussed in the last chapter. Fossil-fuel combustion today releases about 8.5 Gton C per year, and this rate is rising exponentially, driven by population growth and economic growth, rising despite increases in energy and carbon fuel efficiency.

Combining the ocean and tropical deforestation, mankind is releasing carbon to the atmosphere at a rate of about 10 Gton C per year. The atmospheric CO₂ inventory is rising at a rate of about 4 Gton C per year. Where is the other 6 Gton C per year? There are two types of natural sinks for CO₂ that are operating today; one in the oceans, and the other on land.

On land, in addition to the deforestation source for carbon, there appears to be a significant amount of carbon being taken up, called the **terrestrial carbon sink**, probably in places that are not being deforested now, or maybe were cut long ago and are now growing back. The deforestation source can be seen and quantified in satellite images, but the land sink for carbon is fairly difficult to see or measure.

The land surface is taking up a lot of carbon today, but we can't find it.

Compared with the ocean, the distribution of carbon on land is very spotty. Recall from Chapter 7 that most of the carbon on land is in the soil rather than in the trees where we could see it. In soils, the amount of carbon depends on recent land-use history, things like fires, agriculture, or forest regrowth. It is difficult to know precisely how much carbon there is on land because the measurements are so variable; you would have to make a lot of measurements to average out all the noise of natural variations. As a result, it would be possible to increase the amount of carbon on land, a little bit here or there, in a way that would be entirely invisible to direct measurements. This seems to be going on, and so the carbon uptake on land is called the **missing sink**.

The best way to estimate the magnitude of the missing sink, in Gton C per year globally, is to put together the rest of the carbon budget and assign to the terrestrial biosphere whatever is left over. Figuring out where it is going, and why, is trickier. The best approach to actually tracking the CO₂ uptake is to measure CO₂ concentrations as the winds blow across the land surface. If the CO₂ concentration goes down as the air blows over a land area, there must be uptake in that area. There is a network of “CO₂ observatories” around the globe, where precise CO₂ concentration measurements are made daily for analyses such as this. The data are analyzed within atmospheric circulation models that predict wind directions, to try to come up the regional distribution of the land sink. Most of these studies conclude that the missing sink is located in the high-latitude northern hemisphere.

There are several reasons why the land may be willing to take up CO₂ as the atmospheric CO₂ concentration rises. One is that with warming, there is a longer growing season. This has been observed in many climate and botanical records. Also, some tundra areas become amenable to conquest by forests.

Rising CO₂ in the atmosphere may directly encourage plants to grow faster, by a process known as **CO₂ fertilization**. Plants run their photosynthetic machinery inside waxy walls on the surfaces of leaves. Gases are exchanged with the outside atmosphere through adjustable vents in the walls called **stomata**. When the leaf needs CO₂ for photosynthesis, the stomata open. The downside of opening stomata, though, is losing water. So if CO₂ concentrations were higher in the outside atmosphere, plants could get the CO₂ they need without losing as much water.

There is no doubt that CO₂ fertilization is a real effect. CO₂ concentrations in greenhouses are typically higher than in the outside atmosphere, one of the ways that greenhouses are good for plants. However, in the real world, plant growth is very often limited by something else other than water stress or CO₂ limitation, such as fertilizers like nitrogen or phosphorus. Scientists do CO₂ fertilization experiments in natural settings by pumping CO₂ continuously into the air upwind of a grove of trees. When the wind changes, they open a different vent so that the target grove is always downwind from a CO₂ source. These experiments go on for years! What they tend to find is an initial growth spurt from CO₂ fertilization, followed by a leveling off at something like the initial rates. The bottom line is that there is still considerable uncertainty in the strength and importance of the CO₂ fertilization effect for the planet as a whole.

Another process that may affect CO₂ storage on land more and more in the future is the temperature sensitivity of respiration in soils, the process that converts organic carbon back into CO₂. Organic carbon decomposes more quickly as it gets warmer. Think of a ham sandwich, half of which is safely stowed in the refrigerator while the other half sets at room temperature. After a few days, which half will be further along on the pathway to degradation? Because of faster decomposition rates, there is very little organic matter in tropical soils, whereas prodigious amounts of carbon have accumulated in high latitude soils, especially frozen permafrost soils. A warmer world with less permafrost and more areas with tropical climate would hold less carbon in its soils. Soils releasing CO₂ to the atmosphere would be a positive carbon cycle feedback, and it seems like it could be a large one.

The land carbon uptake could saturate or reverse itself in a warming world.

Uptake of fossil fuel CO₂ by the oceans is called the **ocean carbon sink**. Early Earth scientists supposed that the ocean would dominate the CO₂ concentration of the atmosphere, because there is about 50 times more carbon dissolved in the oceans than there is in the air. The ocean covers 70% of the Earth's surface, and the length and width of the ocean are huge compared with its depth, which averages about 4 km. For these reasons, until the 1950s, many scientists thought that the ocean would be rather quick about taking up any extra CO₂ we put in the air.

But the circulation of the ocean tends to limit contact between the deep ocean and the atmosphere. The deep ocean fills up with the densest water, which in our ocean means the coldest waters. Once a water parcel sinks, its temperature does not change much. The deep ocean fills up with water that last saw the atmosphere in very cold places, such as the Antarctic and the North Atlantic, in their local winters. The surface of the ocean is huge, but the deep ocean, which is the largest water type in the ocean, sees only the atmosphere through a very small window in the ocean surface. The deep ocean is so close to the surface, and yet it is quite far away.

Ocean uptake of fossil fuel carbon depends on the circulation. . .

The deep ocean fills up until cold polar waters underlie the warm surface waters in lower latitudes. The warmer waters mix with the cooler, eroding the cold water and making room for more of the coldest water to continue filling the deep sea. As new cold water sinks

from the high-latitude surface ocean into the abyss, it carries with it atmospheric gases like oxygen and fossil-fuel CO_2 , a process known as **ocean ventilation**. It takes many centuries for the waters of the deep ocean to travel through this loop. For this reason, it will take centuries for anthropogenic CO_2 to invade the deep ocean.

Other water masses of the ocean take up fossil fuel CO_2 more quickly but ultimately in smaller amounts. The **thermocline** is a layer of water between the warm surface and the cold abyss, tracing through the water column, a few hundred meters deep, and outcropping in the high latitudes where surface waters are colder. Thermocline waters ventilate to the atmosphere on a time scale of decades. The **surface ocean mixed layer** is formed by turbulence generated by the wind and is typically about a hundred meters deep. This is not as large as the deep sea or the thermocline, but it is a respectable carbon reservoir of its own ventilating CO_2 in about a year.

... and on the buffer chemistry of inorganic carbon in seawater.

When CO_2 dissolves in seawater, it becomes enmeshed in a small but elegant cascade of chemical reactions called **pH reactions**. First CO_2 reacts with water to form **carbonic acid** (H_2CO_3):



Carbonic acid loses a hydrogen ion (that's what acids do, in general, is release hydrogen ions) to form **bicarbonate ion**, HCO_3^-



A second hydrogen ion can be released to form **carbonate ion**, $\text{CO}_3^{=}$ where the superscript “=” means two negative charges.



The concentrations of carbonic acid, bicarbonate, and carbonate ions control the acidity of the ocean, just as they control the acidity of our blood and cell plasma. **Acidity** is a measure of the concentration of hydrogen ions, H^+ , in a water-based solution. Hydrogen ions are very reactive, which is why a strong acid such as the sulfuric acid in a car battery can burn your skin or clothes.

The standard way to describe acidity is using a number called the **pH** of the solution, which can be calculated as

$$\text{pH} = -\log_{10} [\text{H}^+]$$

The $\log_{10}$ is the base-10 logarithm, meaning that if $x = 10^y$, then $\log_{10} x = y$. The hydrogen ion concentration is denoted by the square brackets and is expressed in units of moles of H^+ per liter of solution. A **mole** is simply a specific number of atoms or molecules, called Avagadro's number, equal to $6.023 \cdot 10^{23}$. The hydrogen ion concentration in seawater usually ranges from 10^{-7} to $10^{-8.3}$ moles of H^+ per liter. The pH of seawater therefore ranges from 7 to 8.3.

The pH of a solution describes the balance of H^+ in it. Lots of H^+ is acidic, which is a pH value less than 7. Not much H^+ means lots of OH^- , which is basic, and a pH value greater than 7.


These particular chemical reactions are fast enough that the distribution of chemical species will always be in their lowest-energy configuration, what is called **chemical equilibrium**. Not all

chemical reactions reach equilibrium quickly or even reach it at all, like, for example, the CH_4 and O_2 in the dark bubbles of an ice sheet, which resist the transition to the lower-energy equilibrium state (converting to CO_2 and H_2O). But these pH reactions in seawater reach equilibrium for practical purposes instantaneously.

The response of an equilibrium chemical system, such as the seawater carbon buffer, to a perturbation, such as the addition of new CO_2 , can be described using a simple idea known as **le Chatelier's principle**. The addition of one of the chemicals on, say, the left-hand side of the chemical equilibrium reaction will cause the reaction to shift to the right, the direction that tends to compensate for the perturbation. The principle works in either direction: take some of something out, and the equilibrium will put some of the something back; add more something, and the equilibrium will remove some of the something.

Le Chatelier's principle is treacherous for students of the carbon system in seawater, though, because we have an innate human tendency, I have found, to ignore the hydrogen ions. They are such tiny things, after all. There are far fewer hydrogen ions in seawater than there are dissolved carbon species. What this means, however, is that a small increment of reaction would produce a huge change in the hydrogen ion concentration, relative to how much was already there.

The safest assumption to make is that the carbon species have to get along together without sopping up or sloughing off too many hydrogen ions at all. Hydrogen ion is such a tiny slush fund, it might as well not exist. Reactions 10-1, 10-2, and 10-3 can be combined into a single reaction that does not allow any production or consumption of hydrogen ions—just interconversion of the carbon-containing molecules and water:



To this reaction we can apply le Chatelier's principle with impunity. If we were to add CO_2 to this system, the equilibrium would compensate somewhat by shifting to the right, consuming some of the CO_2 by reacting it with $\text{CO}_3^{=2-}$.

CO_2 dissolves in seawater by reacting with $\text{CO}_3^{=2-}$, an “antacid”, in the carbonate buffer system.

Seawater has the capacity to absorb or release more CO_2 than it would if CO_2 had no pH chemistry, because of the other carbon reservoirs, HCO_3^{-} and $\text{CO}_3^{=2-}$. The chemical system is called a **buffer** because changes in CO_2 concentration are stabilized, or buffered by the shifting equilibrium reaction. Most other gases, like O_2 and N_2 , do not have anything like this buffer chemistry.

If you double the CO_2 concentration of air, a bucket of water in equilibrium with the air will want to roughly double the dissolved concentration of CO_2 in the water. But, because of the buffer system, it takes a lot more CO_2 added to the water to double the CO_2 concentration. Most of the CO_2 you add reacts with $\text{CO}_3^{=2-}$ to make HCO_3^{-} .

It is like sitting at a poker game, with a rich partner standing behind you covering about 90% of your losses and taking about 90% of your winnings. You could endure a losing streak 10 times longer, ultimately losing 10 times as much money, because your partner is buffering your bank. A bucket of seawater can absorb about 10 times more CO_2 than it would if this pH chemistry did not exist.

The amount of CO_2 that the buffer can absorb depends on the concentration of carbonate ion, $\text{CO}_3^{=2-}$. Carbonate ion is like anti- CO_2 . In natural surface seawater, there is about 10 times more $\text{CO}_3^{=2-}$ than dissolved CO_2 , so that the buffer strength is about a factor of 10, the same as in the poker game. But, as more CO_2 is added to the water, the proportions of CO_2 and $\text{CO}_3^{=2-}$

change, CO_2 going up and CO_3^{--} going down, and the buffer gets weaker. This will result in a slowing of ocean uptake of human-released CO_2 in the future.

The relative concentrations of CO_2 and CO_3^{--} in seawater determine its acidity or pH. In an acidic solution like a fizzy cola, the carbon is pretty much all in the form of dissolved CO_2 , which is for this reason called the acidic form of carbon. The opposite of acidic is called **basic**. Make the solution more basic, and more of the carbon will be in the basic form of carbon, which is CO_3^{--} .

Fossil fuel CO_2 makes seawater more acidic, to the detriment of life forms that make shells of CaCO_3 .

When CO_2 is added to seawater, the pH of the water becomes more acidic. This phenomenon is called **ocean acidification**. The surface ocean is already about 0.1 pH units more acidic than it was naturally.

Organisms that make shells out of CaCO_3 will feel the changing pH most strongly. Corals come to mind first (see Chapter 12), but other organisms that make CaCO_3 include coccolithophorids, which are phytoplankton (free-floating microscopic plants) at the base of the ocean food chain, and pteropods, swimming animals with shells like snails that are a primary food source for some salmon, mackerel, herring, and cod.

The pH chemistry of seawater is counterintuitive in that when you put CO_2 in, it causes CaCO_3 to dissolve. One might have thought that adding CO_2 to the oceans would lead to an increase in the amount of carbon that winds up in CaCO_3 , the rising tide floating all boats. It goes the other way because of the change in acidity. CaCO_3 is chemically a base, and if you put it in acid, it dissolves. The carbon from the CO_2 and from the CaCO_3 winds up as the neutral form, HCO_3^- , dissolved in the oceans.

The future of ocean uptake of fossil fuel CO_2 may also be affected by changes in the circulation of the ocean. Surface warming is expected to be most intense in high latitudes because of the ice-albedo feedback (Chapter 7). If the high latitudes warm, the overall circulation of the subsurface ocean may decrease. The circulation of the ocean may stagnate, slowing uptake of CO_2 .

Biology in the ocean acts to decrease the CO_2 concentration of surface waters, by converting CO_2 into organic carbon via photosynthesis (Chapter 8). Dead phytoplankton sink out of the surface waters, carrying their carbon to the deep sea. This process is called the **biological pump**. If all the life in the ocean were killed, that is, if the biological pump were stopped, then the CO_2 concentration of the atmosphere would rise. In the other direction, stimulating the biological pump to work harder could eventually decrease the CO_2 concentration of the atmosphere (see the discussion of ocean fertilization in Chapter 13).

In contrast to the missing sink on land, there are many fairly reliable ways to measure the uptake of CO_2 into the oceans. One way is to measure the dissolved CO_2 concentration in surface ocean water, from which you can calculate how quickly CO_2 is dissolving or evaporating across the sea surface. To get a global CO_2 flux takes lots of measurements (about three million have been made and compiled so far) because the temperature and chemistry of the surface ocean vary a lot from place to place and through the seasons. The net CO_2 invasion into the ocean is found to be about 2 Gton C/year. This is a small imbalance between the much larger CO_2 exchange fluxes between the atmosphere and oceans, which are about 100 Gton C/year. CO_2 goes out of the ocean in some places and into the ocean in others, driven by the huge variability in the temperature and chemistry of the surface ocean.

It is easier to measure carbon uptake into the ocean than into the land surface.

There are computer models of the circulation, biology, and chemistry of the ocean that can be used to estimate human-released CO_2 uptake. The models can be tested and calibrated using measurements of chemical **tracers** in the ocean like carbon-14 and industrial chemicals like Freons in the ocean, which show the pathways by which atmospheric gases flow into the deep ocean. The models agree with the surface CO_2 data, that the oceans are taking up about 2 Gton C/year.

Another way to measure the balance between carbon uptake into the land versus the ocean is to measure the change in atmospheric oxygen that accompanies the rising CO_2 . These measurements can distinguish between land uptake of carbon, which releases O_2 , versus CO_2 dissolving in the ocean, which does not release O_2 . There's a lot more O_2 than CO_2 in the atmosphere, so the changes in O_2 are tiny and have to be measured precisely. The story they tell is consistent with the other methods, about 2 Gton C/year into the oceans.

By difference, it works out that carbon uptake by the missing sink on land is taking up quite a bit of carbon, about 4 Gton C per year. This is more than the deforestation CO_2 release rate of about 1.5 Gton C per year, so the land surface overall is currently taking up about 2.5 Gton C/year, about a quarter of our emissions.

A reconstructed history of the cumulative carbon fluxes is shown in Figure 10-3. The land surface was a net carbon source originally, with deforestation outpacing any natural uptake. Then the missing sink started kicking in just in the last few decades. Cumulatively, the land surface has been a net wash, neither a large source or sink overall. What the future will bring is difficult to tell. Ultimately, there is enough fossil fuel carbon available (5,000 Gton C) to overwhelm any uptake by the land surface (1,500 Gton C including soil carbon). The land surface is important in the carbon budget today and for the next few decades, but ultimately the fossil fuel CO_2 will be cleaned up by the oceans.

Eventually, after a few centuries or millennia, the human-released CO_2 will spread out among the carbon reservoirs of the atmosphere, ocean, and land surface. Computer models of the ocean carbon chemistry, physics, and biology predict that in this new equilibrium state, the atmospheric CO_2 concentration will be higher than it was before the CO_2 was released. The extra CO_2 molecules in the air will not be the exact same molecules that came from fossil fuels because carbon is always passing back and forth among the different reservoirs. But the amount




Figure 10-3 Reconstructions of past CO_2 fluxes due to human activity. From P. Tans, 2009.




Figure 10-4 Long-term fate of fossil fuel CO_2 . Reprinted from Archer, 2005.

of extra carbon in the air represents a significant fraction of the human-released carbon slug, around 25% of it, more or less (Figure 10-4).

After the CO_2 peak comes the long tail, which will continue to affect the climate for 100,000 years.

Some of that extra CO_2 in the air will go away when the pH of the ocean recovers, driven by an ongoing CaCO_3 cycle in the ocean called the **CaCO_3 pH stat**. Dissolved CaCO_3 flows into the ocean in river water, the product of weathering (dissolving) rocks on land. Solid CaCO_3 is re-formed from the dissolved ingredients in the ocean to make shells and coral structures. Most of the shells and corals eventually redissolve, but some fraction of the CaCO_3 that forms every year is buried in ocean sediments. This is the mechanism for CaCO_3 removal from the ocean.

Eventually, the budget for dissolved CaCO_3 into and out of the ocean has to balance, just like the water fluxes eventually balance in the kitchen sink analogy. The way the ocean accomplishes this is by changing its pH. CaCO_3 is a base, and it dissolves in acid. So there would be less burial of CaCO_3 in an acidic ocean. In a high-pH or basic ocean, CaCO_3 would tend to form, like lime crusting out on a faucet.

Speaking of faucets, the CaCO_3 pH stat is another application of the kitchen sink analogy. The water level would be the ocean pH, which depends on the amount of dissolved CaCO_3 in the ocean. The burial rate of CaCO_3 depends on ocean pH the way flow down the drain depends on the water level. A deeper water level in the sink corresponds to a basic ocean, increasing CaCO_3 burial, and a shallower water level would be a more acidic ocean.

Remember that the ocean will be acidified by the invasion of new CO_2 . This would be analogous to scooping some water out of the kitchen sink, slowing the flow of water down the

drain. As part of the oceans' recovery, some CaCO_3 would dissolve on the ocean floor, which in the sink analogy could be mimicked by allowing water to flow back up the drain pipe if the water level in the sink got too low. Also, the warmer climate will probably increase the rate of weathering on land, which would be like a burst of extra flow from the faucet when the water level went down. Both these effects would tend to push the pH of the ocean back toward its natural balance.

As the buffer chemistry recovers, the airborne fraction of the extra CO_2 drops until, after a few thousand years, when the CaCO_3 budget is back in balance, the airborne fraction of the extra CO_2 in carbon cycle models drops to about 10%. What would happen in the real world could be different if there are feedbacks in the carbon cycle, such as the positive feedbacks that appeared to operate through the glacial-interglacial cycles (Chapter 8).

The climate cycle will ultimately recover from the fossil fuel era when the carbon returns to the solid Earth as a result of the silicate weathering CO_2 thermostat from Chapter 8. This mechanism is thought to control the CO_2 concentration of the air on time scales of hundreds of thousands of years. The longevity of the global warming climate event stretches out into time scales of glacial-interglacial cycles, time scales that are longer than the age of human civilization by far. The climate impacts of the long tail are most pronounced in slowly changing components of Earth's climate, such as the ice sheets in Greenland and Antarctica (Chapter 12) that would ultimately drive the largest changes in sea level.

TAKE-HOME POINTS

- The ozone hole problem is not the same as global warming. They are different issues.
- Methane has about a 10-year lifetime in the atmosphere, so its concentration reaches an equilibrium after about this long.
- The land surface and the ocean are absorbing some of our fossil fuel CO_2 , but this could slow or reverse in a changing climate.
- Releasing fossil CO_2 to the atmosphere will affect climate for hundreds of thousands of years—as far as we are concerned, forever.

STUDY QUESTIONS

1. What are the links between the ozone hole and climate change? In what ways are they different from each other?
2. What are the processes that remove anthropogenic methane and CO_2 from the atmosphere? Are those processes reliable, or could they change in the future?
3. Describe the change in the carbonate buffer chemistry in seawater in response to addition of CO_2 . Why does adding CO_2 make seawater more acidic?

FURTHER READING

Archer, D. Fate of fossil-fuel CO_2 in geologic time. *Journal of Geophysical Research Oceans*. doi:10.1029/2004JC002625, 2005.

Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory. Online Trends: A Compendium of Global Change. Available at: <http://cdiac.esd.ornl.gov/trends/trends.htm>.

Intergovernmental Panel on Climate Change (IPCC). The carbon cycle and atmospheric carbon dioxide. In: IPCC Scientific Assessment. Cambridge, England: Cambridge University Press, 2001. Available at: http://www.grida.no/climate/ipcc_tar/.

Tans, P. An accounting of the observed increase in oceanic and atmospheric CO₂ and an outlook for the future. *Oceanography*. 22(4):26–35, 2009.

Weart, S. *The Discovery of Global Warming*. Cambridge, MA: Harvard University Press, 2003. The history book again. Read about early measurements of atmospheric CO₂.

Exercises

1. **Long-Term Fate of Fossil Fuel CO₂.** Use the online geologic carbon cycle model at <http://understandingtheforecast.org/Projects/geocarb.html>. Use the default setup of the model, and notice that the CO₂ weathering rates, etc., for the transient state are the same as for the spin-up state. So, if there were no CO₂ spike at all, there would be no change in anything at year 0. (Go ahead, make sure I am not lying about this.) Release some CO₂ in a transition spike, 1000 Gton or more or less, and see how long it takes for the CO₂ to decrease to a plateau. There are two CO₂ plots in the output, one covering 100,000 years and one covering 2.5 million years. How long does it take for CO₂ to level out after the spike, according to both plots?
2. **Effect of Cutting Carbon Emissions.** Look at the online Integrated Science Assessment Model (ISAM) global warming model at <http://understandingtheforecast.org/Projects/isam.html>.
 - a) Run the model for the “business-as-usual” case (Scenario A); note the pCO₂ concentration of the atmosphere in the year 2100.
 - b) Estimate the decrease in fossil fuel CO₂ emission that would be required to halt the increase in atmospheric CO₂, based on the present-day CO₂ fluxes into the ocean and into the terrestrial biosphere. Test your prediction by imposing these fluxes from the present day onward to the year 2100.
 - c) Repeat experiment 2b but delay the cuts in fossil fuel emissions to the 2050. What is the impact this time on the maximum pCO₂ value we will see in the coming century?
3. **Climate Sensitivity of this model.** Deduce from the above results or new model runs: what is the climate sensitivity, ΔT_{2x} , assumed in this model?

Part III

The Forecast

The human impact on climate can be seen in temperature trends since the 1970s

Is the Globe Warming?

Greenhouse theory is all well and good, but people understandably want to see it in action in the real world before they take it too seriously. They want to see a “smoking gun,” some proof of climate changing as a result of human activity. Greenhouse theory, codified in the numerical climate models, can be tested by comparing the model results against climate changes in the past. Even better is to look for predictions that can be made in advance so that they can afterward be put to the test of whether they came true or not.

The search for a smoking gun.

The first question is how much warming has actually occurred. Answering this question is trickier than one might expect because temperatures vary from place to place and through time, and measurement techniques have changed. One measure of success is to find global temperature reconstructions from independent methods that agree with each other.

The next task is to come up with the histories of natural and anthropogenic **climate forcings**, changes in some parts of Earth’s energy budget that affects the temperature of the Earth. These could include changes in the intensity of the sun, scattering of visible light by human-released and volcanic aerosols in the atmosphere, and changes in greenhouse gas concentrations. Humans affect climate most strongly through greenhouse gas concentrations and aerosols, plus there are a few other, smaller, so-called anthropogenic climate forcings.

Thermometer Records

To see a trend in the global average temperature of the Earth, all the variations—between day and night, through the seasons, and through the weather—have to be averaged out. The temperatures have to be weighted so that parts of the world that have lots of historical weather observations do not dominate the global average, at the expense of places with sparse data. To get a temperature trend across decades of time, we also need to know whether there were changes in the methods used to make the measurements through time. It would be treacherously easy for some bias in the raw data to creep into an apparent change in global temperature.

On land there is a potential source of bias known as the **urban heat island effect**. Sunlight hitting pavement warms it up, the light energy converted to what climatologists call **sensible heat**, “sensible” because you can measure it with a thermometer or feel it with your hand. When sunlight hits land with plants on it, it drives water to evaporate, carrying away **latent heat** (Chapter 5), which you cannot feel or measure with a thermometer until the water condenses. A vegetated landscape does not warm up as much as a paved landscape because some of the heat

is escaping invisibly, or insensibly, I suppose one would say, as water vapor. The urban heat island is an observed phenomenon; cities can be 5°C warmer than their surrounding countryside. Nighttimes in particular are warmer in cities.

What if a weather station were established a century ago in some idyllic rural area, only to have the landscape urbanize since then? The temperature record would show a warming, which would be real for that location, but it would not be due to chemical changes in the atmosphere. Or, if cities were overrepresented in the temperature data, the computed average temperature could be higher than the real average temperature of the Earth.

Independent competing studies find about the same amount of global warming, so their answer seems to me fairly reliable, even though I did not do the analysis personally.

Many independent studies and organizations such as NASA have estimated the changes in the global temperature from the thermometer data, all making some attempt to throw out data from locations that might have seen urbanization (Figure 11-1). The studies all agree that removing urban locations has only a small impact on the global average temperature trend that you calculate from the data. They find that it does not matter how they deal with the urban data; they get more or less the same answer regardless.

The sea surface temperature, or SST, provides an independent temperature record. SST was at one time routinely measured by heaving a bucket on a rope into the ocean to collect surface water, measuring the temperature with a thermometer on deck. Evaporation carried away heat, cooling the water in the bucket somewhat. Some of the buckets were made of wood, others of canvas. Water in canvas buckets cooled more than water in wooden ones because wood is a better insulator. After 1942, mariners began measuring the temperature of surface water as it was pumped into the ship to cool the engine. This provided a more accurate measurement, but there is still the possibility of bias. The waters in the surface ocean sometimes get cooler with depth. A large ship might sample waters from a few meters below the waterline, which might be systematically a bit cooler than the real SST. These effects must be corrected for if




Figure 11-1 Instrumental land temperature reconstruction, from nine independent studies published between 1881 and 2006. Replotted from the Intergovernmental Panel on Climate Change, 2007.




Figure 11-2 Land temperatures, reconstructed from land thermometers, and predicted by using sea surface temperature reconstructions, corrected and uncorrected for the bucket effect. Replotted from the Intergovernmental Panel on Climate Change, 2001.

we are to calculate a global average value of SST through time. The largest correction that had to be applied to any of the temperature records was the bucket correction for the sea surface temperatures. The abrupt change in SST in 1942 resulted from a switch between bucket and engine-room temperature measurement.

The land and sea records cannot be directly compared because the land tends to warm up more than the ocean. The ocean has an infinite supply of water to evaporate and carry away latent heat, and heat can be mixed down from the ocean surface more quickly than heat can penetrate the soils and rocks on land. To take these things into account, the temperature records are compared (Figure 11-2) within a climate model in which the ocean temperatures are set to equal measured temperatures in the real world. Model-predicted land temperatures are then compared with the measurements from land areas. The two independent data sets, each with their own potential biases hopefully corrected for, appear to be consistent with each other.

Sea surface temperatures agree that the Earth surface on average is warming.

Michael Crichton, in his fiction book “State of Fear,” showed temperature records from selected locations around the world, where temperatures have cooled over the decades, a technique called “cherry picking”. These do not disprove the warming trend for the global average temperature, however, because these cooling areas are more than counterbalanced by places that Crichton did not show that are warming, such as the Arctic. The global average temperature record is based on thousands of individual locations around the world. The existence of a few places where climate is cooling does not disprove the warming trend in the global average.

Temperatures Measured from Satellites

Satellites provide another source of global temperature trend information by measuring microwave light emission from oxygen gas, O_2 , in the atmosphere. Because of the way electrons are distributed around the nuclei of oxygen molecules, when oxygen molecules rotate, they emit electromagnetic radiation in the microwave part of the spectrum. This is longer wavelength, lower energy radiation than the infrared (IR) that greenhouse gases absorb and




Figure 11-3 A comparison of the global average temperature reconstruction with the satellite record. Data replotted from Mears and Wentz, 2005.

emit. It is far from the main IR wavelengths of the blackbody spectrum of the Earth (Chapter 4), which is why O_2 is not considered a greenhouse gas like CO_2 . The intensity of microwave radiation emitted from O_2 , however, increases with increasing temperature, just as the intensity of IR goes up with temperature.

Temperature estimates from satellites used to disagree with the thermometers on the ground, but this was caused by errors in data processing that have now been fixed.

These satellite temperature estimates have been the subject of considerable scientific and political discussion because for many years the satellite estimates of warming disagreed with the reconstructions from thermometers. This has now changed, as scientists learn how to calculate near-surface temperatures from the raw satellite sensor data (solid black line in Figure 11-3). The satellites measure microwave emission at several wavelengths and use them to calculate atmospheric temperatures in different altitude ranges. The lowermost altitude range they can resolve spans from the surface to about an 8-km altitude, with an exponential decay kind of a shape, like the curve describing pressure with altitude (Figure 5-2). Even though the weighted average is affected most strongly by temperatures near the ground, there is still some influence from temperatures higher up, including in the stratosphere, where it has been getting colder, masking some of the warming below in the satellite signal until it was corrected for.

The satellite temperature record comes from a series of satellites, which must be well-calibrated against each other as they are pieced together into a longer composite record. The raw data must be corrected for things like changes in the orbit of the satellite. Satellites are ideal for global snapshots, but are not as well-suited to detecting long-term trends. However, the current calibration of the satellite data shows good agreement with the land and sea instrumental records.

Glaciers

The most visual indication of climate change is the melting of mountain glaciers. The vast majority of the mountain glaciers, all around the world, are melting (Figure 11-4). Figure 11-5 shows records of the lengths of 20 glaciers. Many have been melting since the 1700s as a result of natural warming since the Little Ice Age. However, the rate of melting has accelerated in

Qori Kalis



Figure 11-4 The disappearance of the Qori Kalis glacier in the Peruvian Andes.




Figure 11-5 Glaciers are melting all over the world. Replotted from the Intergovernmental Panel on Climate Change, 2001.

recent decades. Mountain glaciers tend to disappear at the altitude where the temperature is somewhat above freezing, so it is fairly straightforward to relate glacier lengths to temperature changes. The glaciers are in general telling the same story as the thermometers, another independent confirmation that the surface of the earth is warming.

Changes in Climate Forcings

Different factors that can change the climate can be compared in terms of their effect on the energy budget of the Earth, in units of W/m^2 , which is called their global average **radiative forcing** values. For example, $1 \text{ W}/\text{m}^2$ in solar output has a similar climate impact as $1 \text{ W}/\text{m}^2$ of IR light trapped by CO_2 , not exactly but they are not too far off. So the radiative forcing makes a convenient starting point for comparing different drivers of climate.

The radiative forcing in W/m^2 allows us to compare the causes of climate change on a roughly equal footing.

The radiative forcing is defined as the energy imbalance immediately after the change has been enacted but before the climate has had a chance to respond. It's defined this way because when the climate responds, it eventually drives the energy imbalance back to zero. Some factors that affect climate are considered internal to the climate system, like the water-vapor feedback, clouds, and changes in the albedo of the land surface and from melting ice. These are not considered forcings, which should be externally driven; because they are internal to the climate system, they are considered feedbacks.

Radiative forcings are related to Earth's temperature by means of a new version of a **climate sensitivity parameter**. The usual way to describe the sensitivity of climate is the equilibrium warming you get from doubling the CO_2 concentration of the air, a parameter written as ΔT_{2x} , described in Chapter 4. The new climate sensitivity relates the equilibrium warming to the radiative forcing; it could be called ΔT_R . A middle-of-the-road estimate for ΔT_{2x} is 3°C , which corresponds to ΔT_R of about $0.75^\circ/(\text{W}/\text{m}^2)$. The current anthropogenic climate forcing today is about $1.5 \text{ W}/\text{m}^2$ (Figure 10.2). Multiplying by the climate sensitivity gives us a temperature change estimate of about 1.1 K. This is more than the Earth has warmed yet because the oceans are keeping us cool (Chapter 12).

We need to know how the climate forcings have changed through time. **Solar intensity** varies over the 11-year sunspot cycle, by about 0.2 to $0.5 \text{ W}/\text{m}^2$ (Figure 11-6). Direct measurements of the intensity of the sun are best made from satellite, avoiding interference from the atmosphere. These measurements go back to about 1975. Earlier than that, the solar intensity can be reconstructed from the number of sunspots that were observed through time, using records that go back to Galileo's invention of the telescope about 400 years ago.

The sun varies in its intensity, through the 11-year sunspot cycle and through the centuries, but it hasn't been getting hotter over the last 50 years, so it cannot explain the warming since then.

Even earlier still, the solar intensity can be estimated from measurements in ice cores of particular radioactive atoms that were produced by cosmic rays hitting the atmosphere, such as ^{10}Be and ^{14}C . These are called **cosmogenic isotopes**. The idea is that a brighter sun has a stronger magnetic field that serves as a better shield against cosmic rays. So a high deposition rate of ^{10}Be at some depth in an ice core is interpreted to mean that the sun was weaker, unable to deflect cosmic rays, and therefore not very bright at that time. It sounds a bit convoluted, but they do in fact observe a correlation between cosmic rays, ^{10}Be production, and solar intensity over the 11-year sunspot cycle in recent times (Figure 11-7).

At times in the past, such as the Maunder Minimum, from about 1650 to 1700, there were no sunspots. This was the coldest period in Europe in the last 1,000 years, and glaciers advanced all around the world, including in the southern hemisphere, so it is fairly clear that the solar luminosity was less than it is now. However, we have no direct solar measurements from a time when the sun acts like this, so we don't have a very solid constraint on how much lower




Figure 11-6 Reconstructed history of radiative climate forcings from solar variability, greenhouse gases, anthropogenic aerosols, and volcanic particle emission to the stratosphere. Replotted from Crowley, 2000.




Figure 11-7 History of sunspot number, and cosmogenic isotope production over the past 1000 years. Data from Beer, 2000, and Lean, 2000.

the solar output was. You can see that there is a factor-of-two uncertainty in the solar forcing in the past, as the thickness of the gray region in Figure 11-6. In addition to the 11-year sunspot cycle, the intensity of the sun drifts up and down on timescales of about 100 years.

Volcanic eruptions inject sulfate aerosols and dust particles into the atmosphere. Particles injected into the troposphere last only a few weeks before they are rained out, but if the eruption is large enough to get particles into the stratosphere, they can spend several years floating around before they eventually sink out. The Mt. Pinatubo eruption in 1991 cooled the planet by 0.5°C for several years. **Volcanic climate forcings** look like randomly spaced spikes of short but intense radiative forcing in Figure 11-7.

To the two natural climate forcings we add the two main anthropogenic ones, changes in **greenhouse gases** and anthropogenic **aerosols**. Neither was important before about the year 1750. Greenhouse gases warm, and sulfate aerosols cool (Chapter 4). The radiative impact of greenhouse gases is fairly uniform globally, whereas the climate forcing from sulfate aerosols is concentrated in the northern hemisphere, downwind from industrialized and heavily populated areas.

The Smoking Gun

A comparison of the historic temperature record with model results is shown in Figure 11-8. The observed temperatures in the heavy black lines rose in the period from 1900 to 1950, declined from 1950 to 1970, and then started rising again from 1970 to today.




Figure 11-8 Model simulations of the climate of the last 140 years using human plus natural natural climate forcing (*top*), and natural forcings only (*bottom*). Repotted from the Intergovernmental Panel on Climate Change, 2007.

The temperatures are compared with results from climate models that participated in an exercise coordinated by the Intergovernmental Panel on Climate Change, or IPCC, for its 2007 Fourth Assessment Report. Nineteen models participated, from around the world, and they are all driven by the same scenarios only. Each model was run twice, once with natural climate forcings only, and once by natural plus human climate forcings. The trajectories of many model runs are plotted on top of each other, looking like tangled clouds of lines, and they were also averaged to make the heavy grey lines in Figure 11-8.

Only the models that were told about human climate forcings were able to simulate the warming since the 1970s—the smoking gun.

The model simulations do a fairly good job of capturing the temperature rise in the early part of the twentieth century, but only the runs that include the human impact on climate are able to reproduce and explain the warming in the last few decades. This plot is considered to be the “smoking gun,” proof that humans have already affected the climate of the Earth.

The smoking gun does not rely on models. Climate scientists have used measured warming and the rise in greenhouse gases, without models, to infer what the climate sensitivity of the real world is. The climate sensitivity, ΔT_{2x} , is the amount of warming that you would get after a few centuries if you doubled atmospheric CO_2 and held it there (Chapter 4). A half-dozen studies, shown in Figure 11-9, were based on temperature changes in the last few decades, and three others were based on reconstructed climate changes from prehistoric times, like the warming at the end of the last glacial period. The studies come up with a most likely range between about 1.5° and 4.5°C . This is a confirmation of the climate models, which predict climate sensitivity values in that same range.

So this is the smoking gun for global warming, the temperature rise since the 1970s. Climate models reproduce the warming when you tell them about rising CO_2 concentrations (and aerosols), and they cannot reproduce the warming when you do not tell the models about the gases. The strength of the argument can be seen by turning things around, asking the question,




Figure 11-9 Probability distributions of the value of the equilibrium climate sensitivity of the Earth, ΔT_{2x} , diagnosed from temperature and greenhouse gas concentration changes through the last decades and reconstructed for the prehistoric past. The studies find it most likely that ΔT_{2x} is between 1.5 and about 4.5°C . From the Intergovernmental Panel on Climate Change, 2007.

What would it take for the global warming forecast to be substantially wrong? Is it possible, or does it seem likely, that humans could put as much CO₂ in the air as we like without changing the Earth's climate, and the whole global warming issue could just go away?

What would it take for the fear of global warming to be wrong?

There are two obstacles to this possibility. First, there would have to be some other mechanism to explain the recent warming since the 1970s: some large perturbation to the energy budget of the Earth that no one has measured or conceived of yet. It takes a lot of energy to warm up the whole surface of the Earth including deep into the ocean, and it cannot happen just by random chaos. Certainly there will always be things about the natural world that are not perfectly understood, maybe even things about which we have no clue yet. But so far, in spite of intensive and highly motivated searching, no one has come up with an alternative source of energy to account for all that warming. This can be seen in the failure of the climate models without the human climate forcings to warm up enough.

Second, there would have to be a reason to throw out the greenhouse effect theory, some explanation of why increasing CO₂ would not block outgoing IR energy by absorbing in its bending vibrational mode, perhaps, or how some negative, compensating feedback might neutralize the warming tendency from the CO₂, somehow allowing the Earth to shed the excess energy without warming up.

Overcoming either of these obstacles seems like a long shot, given how much intense, motivated searching has been done. Overcoming both obstacles together seems extremely unlikely to me. No one would be happier than I would about a reprieve from global warming, but honestly I just can't see it.

The smoking gun warming was forecast in advance.

It's one thing to "hindcast" the past, when you already know the answer, but it's much better to predict the future and then see how it came out afterward. Figure 11-10 shows a global




Figure 11-10 Past predictions of future warming and how they turned out. From the Intergovernmental Panel Climate Change, 2007.

temperature forecast made in the First Assessment Report in 1990, a similar forecast made in 1996 in the second report, and again in 2001 in the third. The actual measured temperatures, and a smoothed trend derived from there, are also plotted for comparison. The first forecast was a bit higher than it turned out to be, the second forecast a bit lower, but in general the agreement is quite good. The Mt. Pinatubo volcanic eruption was also a good test of climate models, which were used to predict the amount of cooling as soon as the eruption was done and they had some idea of how much sulfur went into the stratosphere. Climate models and the real world both cooled by about 0.5°C .

The Climate of the Last Millennium

The thermometer and satellite temperature records indicate a warming trend since about 1970, but this begs the question, Is this warming unusual, or do excursions like this happen all the time? To try to answer this question, paleoclimatologists have extended the temperature record further back in time, before the era of technological measurements. Temperatures from before the days of thermometers can be estimated by a variety of techniques, things that can be measured, interpreted of things we would like to know. These are known as **proxy** measurements.

Proxy temperature measurements will have to do for times before the thermometer was invented.

One source of proxy temperature measurements is the thickness of annual **tree ring** layers in wood (Figure 11-11). In places where winter is much colder than summer, there is a seasonal cycle in the density and color of the wood that grows, leading to the formation of annual rings. Trees grow more quickly during warmer years than during colder years, so it makes sense that tree ring thickness might carry information about temperature. Long records of tree ring thickness can be assembled from collections of many individual trees that lived long ago, by matching the patterns of widths to fit the different trees together in time.

The widths of the rings must be calibrated against temperature changes from part of the record where the temperature is known independently. This is tricky, though, because lots of factors can affect how well trees grow besides the annual average temperature, like the




Figure 11-11 Northern hemisphere mean temperature reconstructions. Replotted from Intergovernmental Panel Climate Change, 2007.

availability of water, nutrients, and even atmospheric CO₂. Young trees grow more vigorously than old ones. The types of trees, and their locations, have to be carefully selected to reduce the impacts of growth factors other than temperature.

Then proxy temperature records from around the world are put together into a global average. When this is done, the “noise” (nontemperature signals) can be reduced still further by developing statistical methods to look for spatial patterns in a map of temperature estimates that match patterns in the climate today, like el Niño, for example. These spatial patterns are used as a kind of filter to extract the temperature signal from the noise in the collected proxy data.

Another proxy measurement for surface temperature changes in the past comes from measurements of the temperature as a function of depth in the Earth and in ice cores, called **borehole temperatures**. The temperature goes up with depth in the Earth or in the ice because heat is flowing from the interior of the Earth to the surface. The more heat there is flowing out of the Earth, the steeper the temperature change with depth will be. If the temperature at the Earth’s surface were held constant for a long time, then the temperature profile with depth in the Earth or the ice would be a straight, linear increase with depth (Figure 11-12).

Now imagine warming the surface of the Earth. The warmth will propagate down into the Earth or the ice, and eventually the whole temperature profile would move to warmer values, with about the same slope as before (assuming the heat flow from the Earth is about the same). At some instant in time, take a snapshot of the evolving temperature profile. Mathematical methods called **inverse methods** can be used to figure out how the temperature at the surface must have changed in the past to generate a given temperature profile snapshot. These borehole temperature records lose the little wiggles in the real temperature record, acting as a low-pass filter on the real temperature changes. Also, the uncertainty gets greater as you try to infer surface temperatures too far back in time. Borehole temperatures show warming throughout the past 500 years, more or less consistent with the tree-ring record (Figure 11-11).

Other sources of information about past climates come from **ice sheets** and mountain glaciers. Glacier lengths have been used to construct a third independent prehistoric temperature reconstruction in Figure 11-11, in substantial agreement with the results from tree rings, boreholes, and, where it exists, the instrumental temperature record.




Figure 11-12 Time evolution of subsurface temperature, when the surface warms up. From a snapshot of the current temperature profile (*black line*), scientists can reconstruct the history of the surface warming (*arrows*).

The proxy records reveal that Earth's climate does vary naturally, driven by fluctuations in the intensity of the sun...

Temperature reconstructions of the last millennium generally show a period of warmth from 800–1200, called the **Medieval Climate Anomaly**, and cooler temperatures from 1350–1800, called the **Little Ice Age**. The Little Ice Age was clearly a global event, as indicated by longer mountain glaciers around the world. The Little Ice Age coincided with a time of low solar flux inferred from cosmogenic isotope production and the **Maunder Minimum** dearth of sunspots (Figure 11-7).

Glaciers do not leave a record from the warmest times, however, and the tree ring records come mostly from the northern hemisphere, so the intensity and footprint of the Medieval Climate Anomaly period are somewhat more difficult to pin down. Unlike today, when the Earth is warming up nearly everywhere, during Medieval time there were regional warmings and coolings through time, but no synchronous “global” warming. Medieval time does not look anything like the changes in our climate since the 1970s.

Models are broadly able to simulate the Medieval Climate Anomaly and Little Ice Age, largely driven by long-term changes in solar forcing. But the climate record of the past millennium is a less stringent test of the climate models than the record from the last century because there is more uncertainty in both the radiative forcing and in the temperature record the models should try to match.

Much has been made of the question of whether the Earth is warmer now than it was naturally during the medieval time. In the end this question is a “red herring”; the smoking gun for global warming is still the thermometer records since the 1970s, not the Medieval Climate Anomaly. Imagine sitting on an electric stove burner in your kitchen. You think the burner is off, but you might have turned the knob when you jumped up there to sit. Should you get off? You feel the current temperature of the burner, and let us say it is 105°F, a bit warmer than it ought to be. But you remember that last summer, during the heat wave, that burner was at 105°F, for natural reasons. Does this mean it is safe now? Not really. A better test would be to pay attention to how quickly the burner is warming up, as an indication of why it is getting warmer this time.

Climate Variations on Orbital Time Scales

As we look further back in time, the uncertainties in climate reconstructions get larger, and the climate changes get more dramatic. The **Last Glacial Maximum**, about 18,000 years ago, was a wildly different climate than today, everywhere in the world; this was especially so in the northern hemisphere, which hosted massive ice sheets, most of which have long since disappeared.

... and, on longer time scales, by variations in Earth's orbit around the sun.

Tree ring and borehole temperatures, used to reconstruct the climate of the last 1,000 years, generally do not work as far back as the Last Glacial Maximum, but constraints on temperatures during the ice age come from sediments in the ocean and lakes and from ice cores. In general, it cooled down more in high latitudes than in the tropics, the “polar amplification” caused largely by the ice albedo feedback described in Chapter 7. Overall, the surface of the Earth was perhaps 5° to 6°C cooler than our natural interglacial climate.

The atmosphere contained less CO₂ and methane at this time (Figure 8.3), resulting in a radiative forcing of about 2 to 3 W/m². All of the ice, in ice sheets, sea ice, and snow, changed

the albedo of the Earth, by about another 3 to 3.5 W/m². Three studies have used the climate of the ice age to diagnose the climate sensitivity, ΔT_{2x} , of the Earth. These are plotted alongside the results from the instrumental record in Figure 11-9, and they mostly get the same answer, something like 3°C for doubling atmospheric CO₂. If the climate models were somehow too sensitive to greenhouse gases, then we would expect the climate models to predict a colder Last Glacial Maximum than that which we piece together from the proxy records.

Past climate changes are consistent with the temperature sensitivity of climate models to changing CO₂.

CO₂ is only part of the reason for the colder glacial climate, but it is striking to look at how similar the CO₂ and temperature records are. Every little wiggle lines up between the two records. It is difficult to argue that CO₂ is somehow unimportant to climate after looking at this record.

The Paleocene Eocene Thermal Maximum

Another climate change episode from the deep past is an interval of time 55 million years ago called the **Paleocene Eocene Thermal Maximum** event. This is much farther in the past than any time we have considered yet, and the data are much scantier.

A deep-past analog for global warming, but so long ago that it is not very useful for testing our understanding.

Measurements of the isotopes of carbon in CaCO₃ shells in ocean sediments tell of a sizable release of carbon to the atmosphere or ocean (Figure 11-13). The carbon that was released had




Figure 11-13 The Paleocene Eocene thermal maximum event. Replotted from Zachos, 2001.

a different ratio of carbon-12 to carbon-13 than the CO_2 in the atmosphere and ocean, and so when it was released it changed the isotopic ratio of the ocean. Changes in the carbon isotope signature of the ocean were recorded in shells of CaCO_3 that formed in that altered ocean and accumulated on the sea floor.

It is impossible to know how much CO_2 was released in the spike because we do not know the isotopic ratio of the carbon in the spike. If it were methane, which has an extremely exotic ratio of carbon-12 to carbon-13, we would need a few thousand Gton C. If it were organic carbon, which has a less extreme carbon isotopic ratio, a larger slug would be required to explain the isotopic excursion. For comparison, the entire inventory of mineable coal, using present-day technology, is about 5,000 Gton C (Chapter 8). The Paleocene Eocene thermal maximum is therefore a reasonable analog to the potential effect of fossil-fuel combustion.

The isotopes of oxygen provide an indication of the warming of the deep sea that resulted from the carbon release. Temperatures were already somewhat warmer in the deep ocean than they are today, before the carbon release. After the carbon was released, temperatures warmed by about 5°C .

It is impossible to use the extent of warming as an indication of the climate sensitivity of the Earth at that time, because it is not known what the CO_2 concentration in the atmosphere was. However, the data serve as confirmation of the prediction, from Chapter 10 and Figure 10-4, that it will take many thousands of years for temperatures to recover to their natural baselines, after a sizable fossil fuel CO_2 release.

TAKE-HOME POINTS

- The past few years have been warmer than the thermometer records of the past 140 years and warmer than reconstructed records over the past 1,000 years.
- Climate change is driven by natural and human-induced changes in radiative forcing, including volcanic eruptions, solar variability, greenhouse gases, and human-released aerosols.
- Models cannot explain the rise in global temperature since about 1970 without including anthropogenic forcing. This is the smoking gun for a human impact on climate.

STUDY QUESTIONS

1. What is a proxy measurement? Describe two types of proxy measurements.
2. What is a climate forcing? Which ones are human in origin, and which ones are natural?
3. What is a radiative forcing, and why was this concept developed?
4. Why is the warming since the 1970s the smoking gun for a human impact on climate?

FURTHER READING

- Bala, G., K. Caldeira, A. Mirin, M. Wickett, C. Delire. Multicentury changes to the global climate and carbon cycle: Results from a coupled climate and carbon cycle model. *Journal of Climate* 18: 4531–4544, 2005.
- Beer, J. Long-term indirect indices of solar variability. *Space Science Reviews* 94: 53–66, 2000.
- Crowley, T. J. Causes of climate change over the past 1000 years. *Science* 289: 270–277, 2000.
- Intergovernmental Panel on Climate Change (IPCC). Scientific Assessment 2001. Cambridge University Press. Available at: http://www.grida.no/climate/ipcc_tar/. Technical Summary.
- Lean, J. L. Short term, direct indices of solar variability. *Space Science Reviews* 94: 39–51, 2000.
- Mears, C. A., F. J. Wentz. The effect of diurnal correction on satellite-derived lower

tropospheric temperature. *Science* 309: 5740–5551, 2005.

Zachos, J., M. Pagani, L. Sloan, E. Thomas, K. Billups. Trends, rhythms, and aberrations in global climate 65 Ma to present. *Science* 292: 686–693, 2001.

The website realclimate.org is a commentary site on climate science by working climate scientists, including myself, for the interested public and journalists.

Exercises

1. Point your Web browser to <http://understandingtheforecast.org/Projects/bala.html>. This is a system to browse the output of a climate model run to the year 2300, from Bala et al., 2005. The browser allows us to explore the output of a coupled atmosphere, ocean, and land surface model. The atmosphere model comes from the National Center for Atmospheric Research in Boulder, Colorado, and the ocean model comes from Los Alamos National Laboratory in New Mexico.

The browser allows several options for making plots.

- a) There are two model runs, one with CO₂ emissions simulating the eventual combustion of the entire coal reservoir, about 5,000 Gton C. This run is labeled “CO₂” on the browser. The other model run is called a **control**, in which there are no CO₂ emissions. This is labeled “Ct1”. You can view results from either of these runs, or you can select a “**Drift-corrected**” option that will subtract changes in the control, the model drift, from the rising CO₂ run.
 - b) Some of what you will see is simple multiyear variability. To remove this, perhaps for comparing one climate state with another, select the **averaging, 10-year climatology** option.
 - c) You can highlight the changes in a climate variable by subtracting one from another, to say things like “it is 2°C warmer than it used to be.” This is called an **anomaly** plot.
 - d) Push the **Plot** button to generate a contour plot.
 - e) Click on a contour plot to display the numerical value from that location on the plot, both for the rising CO₂ run and for the control run. In principle, it should be possible to guess what the value should be using the colors and the contours, but it is often helpful to see a number expressed as a number instead of as a color.
 - f) The **Time Series** option brings up another type of plot entirely, a plot of the year-to-year variations in the global average of the variable. The averaging option does not apply to this type of plot.
 - g) A final option is to view pre-made movies of the model results, using the **Play Movie** button.
2. Use these tools to compare model output with a paleo-temperature reconstruction, say from Figure 11-3.
 - a) The model run begins in year 1871. Bring up this temperature map. Choose some location on the Earth’s surface of interest to you, and that you can find again accurately on the map. The world is your oyster. Click on that location find the temperature there for both the rising-CO₂ and the control model runs. Record these values in a table with columns “year”, “T(CO₂)” and “T(control).” Do this, for the same location, for the next nine years, giving us 20 data points.

- b) Compute the mean of each temperature series using the formula

$$\frac{\sum T}{n},$$

where the capital sigma (Σ) means “take the sum of”, in this case of all 10 temperature values, and n is 10.

- c) Add a pair of new columns to your table, labeled $T - T_{\text{mean}}(\text{CO}_2)$ and $T - T_{\text{mean}}(\text{control})$. Subtract the mean temperatures from the temperature each year, and write this number into each row of these two new columns.
- d) Add another pair of columns and write in the squares of the deviations.
- e) Compute the mean of the squares of the deviations.
- f) Take the square root of the means of the squares of the deviations. This quantity is called the root mean square, or RMS of the deviation, also called the standard deviation, abbreviated as σ . Statistically, about two thirds of the temperatures in the table should be within $\pm 1\sigma$. Is this true in your case? Ninety-five percent of the numbers should be within $\pm 2\sigma$. How many observations ought to be outside of 2σ ? How many do you find?
- g) Move to the year 2000 of the simulation. Start a new table, and record 10 years’ worth of data from the same location as before. Run through the same “rigmarole” to compute the mean and standard deviations of the new temperatures. If the new mean value were the same as the old mean value, you would expect 5%, or 1 in 20, of the data points to be outside of 2σ . How many of the data points actually are outside of 2σ ?
- h) How does the warming you have diagnosed from one location compare with the warming observed in the instrumental record (Figure 11-3)? Is the rising- CO_2 model temperature much different from the control run?
- i) Push the Time Series button to bring up a plot of the global average temperature from the model. The year 1871 is the beginning of the simulation, at model year 250, and the year 2000 is at model year 380. How does the global mean temperature increase compare with your spot estimate?

Drought and storms and sea level, not just the warming

The Warming Itself

The actual global warming part of “global warming,” the change in the average temperature of the Earth, is probably not going to be the most important change in Earth’s climate to affect people on the ground. Most of the social impacts will come from secondary effects like droughts and storms, things that are regional in footprint rather than global. Not even the warming itself will be smoothly spread about the Earth’s surface, but rather some places and times warm more than others. But the average temperature of the Earth is still a good metric, an indicator of the severity of climate change, even if that particular number, the average temperature of the whole planet, is not the actual number that will affect people the most.

Transient climate runs forced by various Intergovernmental Panel on Climate Change (IPCC) emission scenarios predict temperatures 1° to 6°C warmer by the year 2100 (Figure 12-1). The uncertainty in the forecast for the temperature in the year 2100 derives from two sources, which contribute about equally. One is the temperature response to a given amount of CO₂, and the other is how much CO₂ will actually be released.

A temperature change of 1° to 6°C may not sound like much. The daily cycle of temperature is greater than that, to say nothing of the seasonal cycle. The main impacts of future climate change may come from changes in rainfall rather than temperature; but the temperature change by itself is far more significant to the landscape of the world than you might think.

Natural temperature changes of a few degrees centigrade in the past
changed the landscape of the Earth.

One point of comparison is the temperature difference between now and the last ice age, which is estimated to be about 5° to 6°C. This climate change was huge! If we were looking forward to a glacial maximum in the future, rather than backward into the safe past, we would be in a panic. It would be apocalyptic. Europe was an occasionally habitable tundra. The ice sheet in North America dwarfed what exists today in Greenland. Greenland, in spite of its name, is a pretty fierce place. Of course, Europe and North America were extreme cases because the ice sheets were there, but the landscape elsewhere also changed significantly. Traces of changing snowlines make it clear that it was noticeably colder everywhere. Pollen data show huge changes in vegetation type. The coastlines were barely recognizable. Truly it was a different world.

Another comparison is to the Little Ice Age and Medieval Climate Anomaly periods (Chapter 11). Some reconstructions of the global mean temperature, or northern hemisphere temperature, show temperature changes of perhaps 0.5 to 1°C. These climate intervals were not the end of the world, but they definitely sufficed to rearrange civilizations. In medieval time, European agriculture was a bounty of plenty, in a stable, benign climate. Meanwhile the New




Figure 12-1 Model global temperature predictions for the coming century for several different CO₂ emission scenarios labeled B1, A1T, and so on. These are all business-as-usual scenarios with no effort assumed to slow CO₂ emissions for the climate's sake, but they make different assumptions about future trends, along the lines of the Kaya Identity equation in Chapter 9. The gray-shaded areas show the likely ranges of the responses, based on model results (the bars on the right) plus observational estimates of the climate sensitivity.

World was arid, punctuated by strong, decade-long droughts that coincided with the demise of two organized civilizations in the New World, the Classic Maya and the Anasazi.

The climate was much tippier in the Little Ice Age than it was in medieval times. Temperature or rainfall would change suddenly for a year—or for decades. There were periods of drought, of hot summers, of arctic winters, and mild periods of moderate climate and good harvests. In the years 1690–1730, roughly coincident with the **Maunder Minimum** time of no sunspots, there was sea ice around Britain and northern France and a complete rearrangement of the fisheries in the Atlantic.

My impression is that a warming of 1°C would probably not be a world-shattering change, at least globally, although there is the risk of prolonged or perpetual regional droughts, such as occurred in the American Southwest during the medieval warm time. But by analogy to the intensity of the climate changes that came with the end of glacial time, I would guess that a global mean temperature change of 5°C would be catastrophic.

The distribution of the forecast temperature change is not uniform geographically or in time. Plate 12-1 shows the mean annual temperatures from a climate model described in the Exercises section from Bala et al. (2005). The atmosphere part of the model was developed at the National Center for Atmospheric Research, a government agency in Boulder, Colorado, and the ocean part was developed at the Los Alamos National Laboratory in New Mexico. In general, this particular model has a relatively low climate sensitivity, ΔT_{2x} , of 2° to 3°C for doubling CO₂. The temperatures are plotted as anomalies, differences from the temperatures in year 2000 in Plate 12-2.

Global warming is not globally uniform.

The high latitudes warm more than the low latitudes do, by a factor of 3 or 4, mostly because of the **ice-albedo feedback**. Temperatures in Alaska and western Canada have warmed by 3° to 4°C, compared with a global average of perhaps 0.6°C. Models predict that the tundra ecosystem could disappear almost entirely in a business-as-usual world of 2100.

Much of the high-latitude land surface is **permafrost**, defined as soil that is frozen year round. The surface of the soil may melt in the summer, in which case it is called the **active zone**. As temperatures rise, the active zone gets thicker and the soil column can collapse, leaving houses and trees tilted at crazy angles. Most of the Trans-Alaska oil pipeline has its foundation in permafrost soil. Arctic lakes drain away overnight as melting ice leaves holes in the soil column. Coastlines are collapsing at rates of 40 m/year in some parts of the Canadian Arctic and Siberia.

Warming driven by greenhouse gases is generally more intense in the cold parts of the Earth's surface: in wintertime and at night rather than in the summer and in daylight. This is because greenhouse gases impede the loss of heat to space, so temperatures rise just as quickly in the sun but cool down more slowly by shedding of infrared energy. Air over land tends to warm more than air over water because evaporation can always carry away heat from the water, but the land surface may dry out.

An increase in the average temperature of the Earth's surface translates into many more extreme heat events, such as the European heat wave in August 2003, for example, which killed an estimated 35,000 people. It was thought to be a once-in-a-century event, but it was followed by a similar event in 2006. Greenhouse heat waves tend to remain hot all through the night instead of cooling down and allowing people and foliage some cool nighttime reprieve to recover.

Emergent Effects

The rise in global temperature is like a fever in that most of its impact is not directly due to the warmth itself. Changing the average temperature of the Earth scrambles the climate, rearranging weather patterns and water supply in ways that are not easy to forecast reliably and so not easy to prepare for.

Global “weirding.”

Overall, more rain is expected to fall in a warmer world. On the face of it, this sounds like it might be a good thing in a time when freshwater availability is a problem in many places. However, for the rainfall forecast, the devil is in the distribution, even more than it is for temperature (Plate 12-3).

In the tropics, the weather patterns are governed mostly by an overturning pattern called the **Hadley circulation** (Figure 12-2). Air rises convectively at the equator, where the solar heating is most intense. The rising air cools down because of expansion under lowered pressure, and its water vapor condenses into equatorial rain. Air subsides in the subtropics, bringing




Figure 12-2 The atmospheric circulation and the rainfall forecast.

bone-dry air down to the surface and creating deserts at about 30 degrees latitude north and south. Warming the climate is expected to increase the Hadley circulation pattern, intensifying rainfall at the equator and further drying out the desert regions in the subtropics.

Meanwhile, precipitation in the highest latitudes is expected to increase because of the increase in temperature (because there is more water vapor in warm air). This increase in high-latitude precipitation plays a role in the projections for the large ice sheets (see next section).

Rainfall away from the equator can be divided into two types: storminess and monsoons. Storm-driven precipitation is strongest in particular “storm alleys” on Earth (Figure 12-2). The projection is that the rate of storm-driven rainfall would increase in a warmer world because the warmer air carries more water vapor and because the atmosphere would be in general stormier.

Monsoons are driven by a temperature contrast between the land surface and the ocean. The different temperatures drive the winds to blow in ways that systematically dump water as prodigious rainfall. Monsoon systems are sensitive to changes in the temperature differences that drive them, and their strength tends to vary from year to year even in a natural world. According to climate records derived from cave CaCO_3 deposits, monsoons are variable on longer time scales as well, switching on and off abruptly and staying switched for centuries and longer. The projection for the monsoon in southern China and India is to become stronger and stormier. Rainfall in the American Southwest is actually monsoonal in type in that it is driven by Pacific surface temperatures, and it seems to be declining.

Monsoons are particularly vulnerable to climate rearrangement.

Continental interiors are expected to dry out as a result of faster evaporation into warmer air. Even though rainfall globally averaged would increase in a warmer world, areas on Earth that are dry today in general are expected to become drier, increasing the likelihood of drought.

Droughts are driven by changes in the water cycle that vary from one climate model to the next, so forecasts of the impacts in specific regions are not very reliable. This is unfortunate, because drought is probably a more significant danger to agriculture than heat alone. When the results from the climate models from the latest IPCC exercise are combined, however, there is a consistent picture of drying in particular regions, such as the American Southwest, the Mediterranean, and Australia, which are already undergoing drought conditions (Plate 12-4).

Droughts are exacerbated by a feedback from plants. Healthy plants act to capture groundwater, preventing it from running off in a stream but instead evaporating it into the atmosphere, a process called **transpiration**. In this way, plants, and in particular the aggressive plant life of the tropical rain forests, sustain their own existence by recycling rainfall. If the plants die back, rainwater tends to run off, ultimately flowing back into the ocean in rivers rather than re-evaporating into recycled rainfall. The loss of the plants acts to amplify the drought, making the onset of the drought more of a flip-flop than the decline of the rainfall that drove it. The loss of foliage also makes it harder to break out of a drought when rains resume.

Dying vegetation can lock in the drying.

The carrying capacity of the land ultimately depends on water availability. During the Dust Bowl drought in the American Southwest in the 1930s, 85% of the residents of the state of Oklahoma picked up and left. There are proxy records in lake and ocean sediments of droughts in that region in the prehistoric past that were much stronger and more persistent than the Dust Bowl. The Classic Mayan civilization collapsed in several stages that correlated with strong decade-long droughts as documented in lake sediments and cave deposits.




Figure 12-3 The frequency of different intensities of hurricanes through the last decades. From Webster et al., 2005.

Glaciers are expected to melt with warming, and indeed they have already begun doing so (Figure 11-4). This water contributes to a rise in sea level (see next section). Perhaps the most pressing social concern about the loss of glaciers is that they supply water through the summer to a lot of people (see Human Impacts section).

A warmer world would be a stormier world because storms are driven primarily by the heat of water vapor condensing into precipitation. Warmer air carries more water vapor, so there is more energy to be “wrung out” when warmer air rises and cools. Storms and weather are already observed to be getting stronger and more extreme. There have been more hard rains, longer and hotter heat waves, stronger tropical cyclones, and more droughts.

Tropical cyclones, also called typhoons or hurricanes, may get more intense or more prevalent in a warmer world, although there is considerable uncertainty about this. The number of hurricanes per year has been holding steady through the decades, but the frequency of the most powerful hurricanes, Categories 4 and 5 storms, has doubled in the past 30 years (Figure 12-3).

A warmer world would be stormier.

Every year there are about 90 tropical storms, 40 of which become cyclones. These storms grow from particular combinations of winds and pressures that serve as “seeds” for the storms. Whether the seeds grow or dissipate is determined by the structure of winds in the air, whether they blow in the same direction through the entire atmospheric column, or whether they shear the growing storm apart by blowing in different directions at different altitudes.

The hurricane draws its power from the water vapor in the warm air over ocean. Hurricanes do not form if the sea surface temperature is cooler than about 26°C, so their range would expand in a warmer world. Also, everything else being equal, warmer sea surface temperatures would make hurricanes stronger. But the future of hurricanes is difficult to predict reliably because their birth and growth are affected by subtleties in the structure of the atmosphere and ocean, which might change in a warmer world. The strength of a cyclone also depends on the subsurface temperatures of the ocean, because the wind tends to mix cold water up from below, and on the air temperature near the top of the troposphere. The bottom line is, there are good reasons to worry that cyclones might get more intense with rising CO₂, but the connection is not as simple as, say, the link between CO₂ and global mean temperature.




Figure 12-4 The distribution of wind speeds in model hurricanes for a control world (natural climate) and doubled CO₂. From Knutson and Tuleya, 2004.

Model-simulated hurricanes are more intense in a doubled-CO₂ world. Ordinary climate models do not have enough resolution, meaning that their grid points are not close enough together, to simulate hurricanes very well. Hurricanes can be simulated in global models, however, by a technique called **adaptive grid refinement**, essentially inserting grid points into a model when and where they are needed. Alternatively you can think of it as running a second, high-resolution model in the area where a hurricane is, coupled to and exchanging information with a low-resolution global model. The model results in Figure 12-4 predict that there will be more category 5 hurricanes in a high-CO₂ world.

A hurricane generates a rise in sea level called a **storm surge**. The low air pressure in the center of the hurricane acts to lift the ocean surface, and the winds tend to blow the water into the center in a pile. The storm surge is the aspect of hurricane Katrina that destroyed parts of New Orleans. If storms get stronger, their higher storm surges will add to the impact of higher sea level, discussed next.

Ice and Sea Level

There are two reasons why sea level would rise in a warming world. One is the **thermal expansion** of the water. Warmer water expands, filling up the ocean basin like mercury rising in a thermometer. Depending on the amount of warming and the circulation and mixing in the ocean, the IPCC forecast calls for about 0.1 to 0.4 m of sea-level rise attributable to thermal expansion alone by the year 2100. This is just the beginning, however, because it takes centuries for heat to invade the oceans. Coupled atmosphere-ocean climate models suggest that a long-term CO₂ doubling might eventually, after a thousand years or so, raise sea level by 0.5 to 2.0 m as a result of expansion.

The other component of sea level rise is from **melting ice on land**. There are several functionally different types of ice in the Earth's climate system. **Sea ice** forms from seawater freezing at the surface and may be augmented by snow that lands on the ice. **Glaciers** are formed from compacted snow on a mountaintop and flow down mountains like silent rivers of ice (Figure 11-4). **Ice sheets** are larger than glaciers and are found today in Greenland and Antarctica. **Ice shelves** are pieces of ice sheets that have flowed out onto water, floating but still attached to the ice sheet.

Sea level changes slowly and will continue to rise for centuries.

Floating ice, the sea ice and ice shelves, does not raise sea level when it melts because the ice already displaces its own weight in water. When floating ice melts, its water exactly fills the hole that the ice had previously occupied. This is **Archimedes' principle**.

Ice on land however does increase the sea level when it melts. Of the land ice, the ice sheets in Greenland and Antarctica hold most of the water by far. Together, the two ice sheets could raise sea level by about 70 m if they melted, about the height of, say, a midrise 14-story hotel on Miami Beach. The mountain glaciers hold much less ice, about 0.5 m of sea level, but because they are smaller, they can respond much more quickly to changes in climate. Today, mountain glaciers are contributing more to the rise in sea level than the major ice sheets are, and about as much as from thermal expansion.

The ice sheets in Greenland and Antarctica are the slow giants of the rise in sea level.

The ice in an ice sheet flows, very slowly, from the center, where it accumulates as snow, to the fringes where it melts or runs into the ocean and breaks up into icebergs. The air at the base of the **Greenland ice sheet** is close to the melting temperature, and ice sheet models predict that the Greenland ice sheet would eventually melt if the world were about 2°C warmer than our natural pre-industrial climate (often defined as the climate of the 1950s). Melting Greenland would raise sea level by about 7 m, the height of a two-story building.

The **Antarctic ice sheet** is colder than the Greenland ice sheet, and most of the ice in Antarctica is thought to be relatively stable, not as sensitive to changes in temperature as that in Greenland. The vulnerable part of Antarctica is the **West Antarctic Ice Sheet**, which holds enough water to raise sea level by about 5 m. This ice sheet is unique in that it is sitting on bedrock below sea level (Figure 12-5). When the ice sheet first started forming, the bedrock was above sea level, but the crust gradually sank under the weight of the accumulating ice. The crust sank most deeply in the center, so the ice sheet is setting in a bowl below sea level.

The concern is that a warmer ocean might drive the edges of the ice sheet to melt, allowing it to float free of the bedrock. Floating ice flows much more easily than grounded ice, pulling more ice from the ice sheet. Because the water gets deeper toward the center of the ice sheet, it would become progressively easier for the ice to float free, leading to an accelerating cycle of float, flow, and melt.




Figure 12-5 A cross-section through the West Antarctic Ice Sheet. Redrawn from Oppenheimer, 1998.

There is geological evidence that the size of the West Antarctic ice sheet may have been smaller during the warmer-than-today Pliocene time, about three to five million years ago. This would suggest that the West Antarctic Ice Sheet might become unstable after 3°C of warming, a slightly higher trigger than for Greenland.

Overall, the IPCC forecast for the rest of the century is for sea level to rise 0.1 to 0.5 m, approximately half attributable to expansion and half to the melting of ice, mostly from small mountain glaciers. Models of Antarctica and Greenland tend to predict that it would take thousands of years for these ice sheets to respond fully and that not much will happen in the coming century. However, recent events, and reconstructions of prehistoric events, give reason to believe that real ice sheets might respond to climate change in ways that the ice sheet models cannot simulate yet. For this reason, the IPCC specifically did not include the possibility of changes in the flow of ice from the sheets into the ocean, what they called “dynamical changes” in the ice, in their latest sea level rise forecast. The forecast comes with a potentially very large asterisk.

Ice on Earth knows a few things about how to melt
that scientists have not figured out yet.

Ice shelves have demonstrated an ability to self-destruct catastrophically. A piece of the **Larsen ice shelf** on the Antarctic peninsula, the size of Rhode Island, collapsed into a sea of icebergs within a few months in 2002 (Figure 12-6). Ice shelves do not themselves affect sea level when they collapse into bergs and melt because they were already floating in the water. However, their presence seems to restrict ice flow from the ice sheet farther inland toward the sea.

Within the ice sheet are rivers of ice called **ice streams**, flowing at speeds of hundreds of meters per year. Upstream of the Larsen B ice shelf on the Antarctic Peninsula, and the collapsing Jakobshavn ice shelf in Greenland, the ice streams accelerated markedly after the ice shelves collapsed. Several of the ice streams draining the West Antarctic ice sheet run into the Ross ice shelf, which is within a few degrees centigrade of the Larsen ice shelf temperature before it exploded.

The significance of iceberg discharge is that it is an extremely efficient method of melting ice. Melting of an ice sheet in place is slow because it takes a long time for the atmosphere to carry that much heat up there. Icebergs move the ice mountain to the sun.

There is also evidence from the prehistoric past of fast collapses of ice sheets. The Laurentide ice sheet, perched on North America during glacial time, purged large amounts of ice into icebergs periodically in what are called **Heinrich events**. The primary evidence is the finding of layers of rocks in open-ocean sediments large enough that they could only have been carried there in floating ice. The sediments tell us that the events lasted a few centuries, and there must




Figure 12-6 Satellite images of the Larsen ice shelf on the Antarctic Peninsula. A huge area of ice shelf exploded into icebergs in just a few weeks.

have been a veritable flood of icebergs, enough to change the circulation of the North Atlantic and alter the climate around the world. The events lasted a few centuries and are estimated to have raised sea level by something like 5 m.

Currently there are no computer models of ice sheets that can reproduce the Heinrich events. There is too much ice-rafted debris spread out over too much of the North Atlantic. Ice sheet models do not collapse that quickly or pick up that much debris. The question is whether the Greenland or West Antarctic ice sheets could also collapse in a similar way. Presumably, if it started, it would be completely unstoppable, running its course over a few centuries. Maybe sea level rise from one collapsing ice sheet could provoke the other to go.

Sea level at any given location on the Earth depends on the total volume of water in the ocean, of course, but also on local geological movement of the ground, up or down. The largest vertical motions on the surface of the Earth today are the result of melting of ice sheets from glacial time. Remove all that weight and the local crust will float higher relative to the surrounding crust. There is also vertical motion in other parts of the world that is comparable to the projected rate of the rise in sea level. Bangladesh is a hard example; the local sinking or subsidence rate of the land there is about 0.5 cm/year, comparable to the projected rate of sea level rise, doubling the inundation that Bangladesh would experience just from the rising ocean.

It might be possible to defend some of the land against the invasion of the ocean. Much of Holland is located below sea level, and the ocean is kept at bay with dikes and pumps. This could presumably work for cities like Manhattan and New Orleans, but long low-lying agricultural areas such as in Bangladesh and Egypt will be more challenging to save. An increase in sea level of 1 m would inundate about 10% of the farmland in these countries, displacing millions of subsistence farmers. Today the Pacific island of Tuvalu imports food because they can no longer grow their own as a result of salt encroachment in their ground water.

The long-term geologic record tells a much more alarming story about Earth's temperature and sea level than one would infer from the IPCC forecast to the year 2100 (Figure 12-7). During glacial time, Earth's average temperature was about 6°C colder than today and sea level was 120 m lower. The Pliocene, three million years ago, comes before the time of glacial-interglacial cycles, a warmer time with sea level 25 m or so higher than today.

During times of hothouse climate, say 40 million years ago, the temperature of the Earth was perhaps 4° to 5°C warmer than it is today. The sea level 40 million years ago is difficult




Figure 12-7 The relationship between sea level and temperature on geologic time scales. The Intergovernmental Panel on Climate Change forecast for the year 2100 is for less than 1 m of sea-level rise, whereas the geological reconstructions suggest that if the Earth stayed 3°C warmer for long enough, sea level could rise by around 50 m.

to pin down because on this long time scale, the continental crust bobs up and down, moving the beaches and corals that we find today in ways that have nothing to do with global sea level. However, the hothouse world was a time when there were no ice sheets on Earth, which would raise sea level on Earth today by 70 m if the Earth returned to that climate state.

Eventually sea level could change 100 times more than the forecast for the 2100.

Sea level in the geologic past changed by 10 or 20 m for each 1°C of temperature change. This is about 100 times more change in sea level than the IPCC forecast for the year 2100. The reason for the difference is the presumption that it will take longer than 90 years for the ice sheets to respond to the change in climate.

As explained in Chapter 10 (Figure 10-4), carbon mined from the Earth and emitted as CO_2 to the atmosphere affects climate for tens of thousands of years into the future, plenty of time for the ice sheets to respond. So ultimately, although ice sheets may respond slowly to climate, the carbon cycle is patient and can afford to wait. The target year of 2100 for IPCC forecasting is set by our own lifetimes and the time scale for some kind of conceivable political decision making. But 2100 is not really the appropriate time scale to discuss the phenomenon of the ice sheets responding to CO_2 . It is as if we saw someone standing at the edge of a cliff contemplating jumping, and we advise that person that in the first 0.1 seconds after he jumped, nothing bad would happen (because the person had not hit bottom yet). We would be using an inappropriate time scale to describe the situation.

Abrupt Climate Change

An abrupt climate shift would be a surprise, like the ozone hole. Paleoclimate records show sudden flip-flops, where the climate changed in a few years and then remained in its changed state for centuries.

The ice core record from Greenland shows a period of stability called the **Holocene**, the past 10,000 years when civilization and agriculture developed (Figure 12-8). Before this, during the last ice age 20 to 80 thousand years ago, the climate of Greenland was less stable.




Figure 12-8 Temperature reconstruction from Greenland. From the Greenland Ice Sheet Project (GISP) II ice core.

Temperatures underwent 1,000-year swings called **Dansgaard-Oeschger (D-O) events**, sudden warmings of 10° to 15°C within a few years, followed in general by more gradual coolings. The D-O cycle is punctuated by the Heinrich events, times when the Laurentide ice sheet collapsed into bergs in the ocean. The climate impacts of these abrupt climate changes were most intense in the high northern latitudes but were felt as far away as New Zealand and Antarctica.

In a way, the IPCC forecasts represent best-case scenarios because they are smooth and there are no surprises.

Scientists are still trying to figure out the reasons for all of the variability in the glacial climate (Figure 12-9). The circulation in the North Atlantic is one player. The Atlantic carries a




Figure 12-9 How the overturning circulation in the North Atlantic might affect the climate.

considerable amount of heat into the high-latitude northern hemisphere. Warm surface waters in the high latitudes cool and sink, making way for more warm surface flow, bringing in more heat. Tracers for the circulation and chemistry of the ocean, preserved in sediment cores, show that many climate changes in the past, such as the Heinrich events, were correlated with changes in Atlantic circulation.

Another ingredient in the climate shifts was sea ice. Cutting heat transport, perhaps capping the ocean with fresh water that would freeze before it would sink, resulted in an expansion of sea ice cover. Ice reflects sunlight, cooling the area further. Air over sea ice can also get much colder than air over open water can. The climate of Greenland, where ice core records come from, is sensitive to massive sea ice formation because it loses its maritime influence.

The abrupt climate change scenario for the future that people have fretted about is a cousin to the cold configuration of the glacial climate, in which the overturning circulation in the Atlantic slows down, allowing the northern high latitudes to cool. This “Day After Tomorrow” scenario was already mentioned in Chapter 7. The overturning could be impeded by an increase in freshwater input to the North Atlantic surface ocean. The largest potential source of fresh water is the Greenland ice sheet, but it would also rain and snow more in the high latitudes of a warmer world.

Current thinking is that the “Day After Tomorrow” scenario, of a return to an ice age triggered, paradoxically, by global warming, is unlikely because the warming influence of the rising CO₂ will be stronger than the local cooling influence of changing the ocean circulation. My personal intuition is that looking too slavishly at past abrupt climate changes to predict the future is like a general in an army always planning for the last war. My bet is that future abrupt climate changes are going to be surprises.

Biological Impacts

The character of the land surface, whether it is hardwood forest or grassland, is determined to a large extent by the local climate. Temperature is important, of course, although it may be the coldest temperature of the year that really matters, or the warmest, rather than the annual average. Precipitation also plays a role in shaping the natural landscape. And of course the human footprint on the land surface includes agriculture, decorative lawns, and pavement.

The infinite complexity of the real landscape can be modeled by simplifying it into some manageable number of different biome types, such as “tropical deciduous” or “savanna” (Plate 12-5). Biome models are based on the physiologies of plants of various types and keep track of plants, metabolic energy, water stress, nutrient availability, and soil properties. Different biomes have different capabilities and tolerances, and they compete with each other for land surface area.

The equilibrium biome response to a new climate can probably be predicted fairly confidently because the distribution of biomes on Earth today is fairly well correlated with climate. The simplest model for terrestrial biomes could simply be a look-up table of the ranges of conditions inhabited by each biome type.

It is more difficult to forecast transitions between one climate state and another. Some organisms, like insects, move easily, but forests have to move by scattering seeds and a slow process called **ecological succession**. When a forest grows after a clear-cut, the first set of plant and tree species will be different from the species that you would find in an old-growth forest under the same climatic conditions. The process of ecological succession is not trivial to understand in the natural world, which makes it even more difficult to forecast for the future.

The biotic response of a climate model is shown in Plate 12-5. The most striking change in Earth’s landscape is in the Arctic. The biome response is strong there because the climate change is intense. The melting permafrost described above also changes the biotic landscape. The model in Plate 12-5 experienced a near total loss of the tundra biome by the year 2300 and the loss of most of it by 2100.




Figure 12-10 Coral bleaching events in Tahiti. Bleaching events occur whenever temperature exceeds a threshold value. A future warming of 2.5°C would be very large compared with the coral stress level. From Hoegh-Guldberg, 2005.

Polar bears are vulnerable to melting sea ice. Polar bears eat only in the winter by fishing for seals through holes in the ice. In summer, when the ice near the coast melts, the polar bears have to live on their accumulated fat. The polar bear as a species may find refuge in a remnant area of sea ice north of Greenland, but where the ice is gone the polar bears will go locally extinct.

In the ocean, coral reefs appear to be vulnerable. Reefs are built by sedentary animals called corals. The corals house symbiotic algae, plants that can perform photosynthesis to aid the nutrition of the corals. Corals almost everywhere in the world are under attack from runoff pollution, clouding the water and cutting their algae off from sunlight. Overfishing threatens corals, as do diseases, invasive species, and acidification of seawater by invasion of fossil fuel CO₂ (Chapter 10). In addition to all these struggles, corals are vulnerable to increases in temperature.

When corals are stressed, they respond by expelling their symbiotic algae. This is called **coral bleaching** because of their loss of color. Bleaching may be a mechanism for the coral to try to find new symbiotic algae that are more suited to the conditions they find themselves in, but it is a measure of desperation. Bleaching is often followed by the death of the coral. Figure 12-10 shows a record of temperatures in Tahiti, with arrows indicating times of coral bleaching events. The correlation between heat spikes and bleaching is very clear. Sediment cores in Caribbean reefs indicate that the degradation of reef communities is more widespread than has been seen in centuries. Corals are also challenged by the acidification of the surface waters from rising CO₂.

In general, the diversity of life (number of species) on Earth is already decreasing because the natural world is restricted and fragmented by human land use. Climate change can only amplify this extinction trend by demanding that natural ecosystems get up and move, just when they are restricted by human land use and therefore least able to do so.

Human Impacts

The strongest direct effect of a warming climate on human welfare will probably be due to changes in water availability. The world is drier today than the more natural world of 1950 and would continue drying in a warmer world. A drought of a few years' duration, or a decade,

can be endured, but rearrangements of the climate system in a warmer world could persist for centuries. Adding to the problem would be melting mountain glaciers and snow pack, which today serve as a source of summertime fresh water in many parts of the world, such as along the Himalayas in India, Pakistan, and Uzbekistan; along the Andes in Lima, La Pas, and Quito; and in the North American Pacific Northwest.

Other aspects of climate changes may also affect crop yields. Corn is sensitive to heat spikes, and just a few can impact the bottom line for the entire season. On the other hand, increasing CO₂ concentration in the air is expected to boost agricultural production by the CO₂ fertilization effect (Chapter 10). Higher CO₂ levels help plants deal with water stress by allowing them to get the CO₂ they need for photosynthesis without opening their stomata as much. Another important factor in the forecast is adaptation. If corn no longer grows in a region because of frequent heat waves, the farmer will plant some other crop that is more heat tolerant.

The story of agriculture in the past century has been explosive, with production growing even faster than population and at ever diminishing costs. The great success of the agricultural revolution seems to argue against worrying about a collapse in global food production.

On the other hand, agricultural yields and the number of people that the Earth can support are products of the infrastructure of our civilization: fertilizers, transportation, and processing of the food. If that infrastructure broke down, it could lead to a collapse of the human population. It has happened before, to the classic Mayans, for example. Sustained regional drought or coastal flooding from a storm surge with rising sea level could lead to distressed migration of millions of people, which would stress the ability of civilized governance to cope. How many of us would survive if there were no food delivered to the grocery stores?

There is a concern that rising temperatures will widen the area where tropical diseases such as malaria and dengue fever will spread. Both have a significant impact on public health, and both appear to be spreading with an increase in the tropical zone of the Earth. In my opinion, the spread of tropical diseases is a serious threat, but it is not the first and most important reason to combat global warming.

In general, the impacts of a changing climate are expected to become stronger and more clearly negative as climate warms more intensely. The impacts would be stronger in the tropics and in the developing world, where there is less wealth to buffer setbacks from the weather and less infrastructure to deal with the aftermath.

TAKE-HOME POINTS

- Climate models that simulate the past predict huge climate changes in the future, relative to what civilized humanity has ever seen before.
- Extreme weather events are already getting stronger and this trend is expected to continue.
- Drought is one of the most dangerous aspects of climate change, as continental interiors warm and dry out, despite a global increase in the amount of rainfall overall.
- Sea level will continue to rise for centuries, from thermal expansion and melting land ice.

STUDY QUESTIONS

1. How can the forecast call for increased rainfall globally and also increased drought?
2. What is the role of vegetation in a drought?
3. Why would storminess increase in a warming world? What kinds of storms could be affected?
4. What are the factors that would increase the risk of storm-surge flooding associated with tropical storms?
5. Why do some kinds of ice on Earth change the sea level when they melt and other kinds do not?
6. Why does sea level take so long to respond to climate?
7. How do reconstructions of sea level changes in the geologic past compare with the forecast for the year 2100? Why?

FURTHER READING

- Archer, D. E., B. Buffett. Time-dependent response of the global ocean clathrate reservoir to climatic and anthropogenic forcing. *Geochemistry Geophysics Geosystems*. 6(3): doi:10.1029/2004GC000854, 2005.
- Bala, G., K. Caldeira, A. Mirin, M. Wickett, C. Delira. Multicentury changes to the global climate and carbon cycle: results from a coupled climate and carbon cycle model. *Journal of Climate*. 18: 4531–4544, 2005.
- Cox, P. M., R. A. Betts, C. D. Jones, S. A. Spall, I. J. Totterdell. Acceleration of global warming due to carbon-cycle feedbacks in a coupled climate model. *Nature*, 408: 184–187, 2000. Intergovernmental Panel on Climate Change, Scientific Assessment, 2001, from Cambridge University Press, or downloadable from http://www.grida.no/climate/ipcc_tar/. Technical summary.
- Emanuel, K. *Divine Wind*. Oxford, United Kingdom: Oxford University Press, 2005. This book is about hurricanes. This is an elegant treatment of the history and physics of hurricanes written by the guy who, well, wrote the book on hurricanes. The book is replete with color photos, plots, and paintings, with a few equations and many quotes from antiquity.
- Emanuel, K. Increasing destructiveness of tropical cyclones over the past 30 years. *Nature*, 436(7051): 686–688, 2005.
- Hoegh-Guldberg, O. Low coral cover in a high-CO₂ world. *Journal of Geophysical Research-Oceans*, 110(C9), 2005.
- Kolbert, E. *Field Notes from a Catastrophe*. New York, USA: Bloomsbury USA, 2006.
- Knutson, T. R., R. E. Tuleya. Impact of CO₂-induced warming on simulated hurricane intensity and precipitation: sensitivity to the choice of climate model and convective parameterization. *Journal of Climate*, 17(18): 3477–3495, 2004.
- Lynas, M. *Six Degrees, Our Future on a Hotter Planet*. Washington, D.C.: National Geographic, 2008. This is a street-level description of what climate change might be like.
- Oppenheimer, M. Global warming and the stability of the West Antarctic Ice Sheet. *Nature*, 393(6683): 325–332, 1998.
- Solomon, S., G. K. Plattner, R. Knutti, P. Friedlingstein. Irreversible climate change due to carbon dioxide emissions. *Proceedings of the National Academy of Science*, 106: 1704–1709, 2009.
- Webster, P. J., G. J. Holland, J. A. Curry, H. R. Chang. Changes in tropical cyclone number, duration, and intensity in a warming environment. *Science*, 309(5742): 1844–1846, 2005.

Exercises

1. Point your web browser at <http://understandingtheforecast.org/Projects/bala.html>. The physics part of the model—temperatures, winds, and currents and all that—was described in the Projects section of the last chapter. In addition, the run had a land biosphere model, developed at the University of Wisconsin in Madison. Plants on land are described by 16 different biome types, which compete with each other for the land surface. The biomes are listed in Plate 12-5. The run also used a simple groundwater scheme, much simpler than real soils. Instead of allowing groundwater to flow as happens in the real world, model soils collect water and allow it to evaporate, like a bucket. If too much water accumulates, it will overflow into rivers and to the oceans.

You can see annual mean surface temperatures and mean surface temperatures for December through February and for June through August. You can also see precipitation, soil moisture content, and vegetation type.

Use the different data-processing and visualization tools described in the Chapter 11 Exercises section to investigate the following questions.

- a) What is the global average temperature increase from the beginning to the end of the simulation? Most climate simulations end at the year 2100, but this one goes to 2300. You may investigate either year as an “end of simulation.”
- b) What is the predicted warming in the winter high latitude, the summer high latitude, and the tropics?
- c) Can you see evidence for poleward migration of biome types? How far do the biomes move?
- d) Can you see a systematic change in precipitation or soil moisture with climate change?




Plate 12-1 Maps of climate model annual mean temperature, from the CCM3 Browser at the book's website. Data from Bala et al., 2005.

Plate 12-2 Maps of climate model temperature changes from the year 2000, from the CCM3 Browser at the book's website. Data from Bala et al., 2005.




Plate 12-3 Maps of climate model precipitation in the year 2000 and precipitation changes from the year 2000. From the CCM3 Browser at the book's website. Data from Bala et al., 2005.




Plate 12-4 Precipitation changes predicted by the Intergovernmental Panel on Climate Change AR4 suite of 22 climate models from around the world.

Source: Solomon, et al. Irreversible climate change due to carbon dioxide emissions. PNAS 2009 106 (6) 1704–1709. Figure 3.




Plate 12-5 Maps of climate model biome types, from the CCM3 browser at the book's website. Data from Bala et al., 2005.

It should not be technically or economically too difficult, but it will be hard to make the decision

Global Warming Is a Large-Scale Problem

Global warming is a difficult problem to cope with politically because its footprint is so large. It is a global problem, subject to an effect known as the **Tragedy of the Commons**. The archetypal example of this effect is of a field for grazing sheep used in common by all residents of a village. Let us say that each new sheep added to the flock decreases the harvestable weight of each of the other cute little sheep by 10% (Figure 13-1). A farmer with one sheep can add a second, doubling his number of sheep. But the result is that all sheep, including the farmer's, lose 10% of their harvestable weight (denoted as “lamb chops” in Figure 13-1). Each will provide 0.9 times as many lamb chops as the original sheep did, so the farmer will now own the equivalent of 1.8 of the original-sized sheep. Our farmer is better off than when he only had one sheep, but all the rest of the farmers lose 10% of their prior sheep output.

In this particular example, there is an overall optimum number of sheep to graze on the common, which will result in the greatest overall number of lamb chops. That optimum number of sheep turns out to be about six. To compute the self-interest of our individual farmer now, let us assume he has one of the six already. The farmer's sheep is worth 0.5 lamb chops at this sheep density. Our farmer buys one more, and in the end has two sheep that will each produce 0.4 lamb chops, for a total of 0.8 lamb chops. It turns out to be in the farmer's self-interest to add sheep, up until there are nine sheep in the yard.




Figure 13-1 The Tragedy of the Commons demonstrated in a totally made-up story about sheep.

Competing interests using a shared resource will tend to overexploit the resource.

The point is that each farmer's personal best interest is served by grazing as many sheep as he or she can afford without much regard to the damage to the field. This results in a tendency to overgraze. From each farmer's point of view, the other villagers are all also overgrazing the common yard, so his own personal sacrifice, to preserve the common, would be in vain anyhow and would put his family at a disadvantage. The end result is far more sheep than can be most efficiently supported on the common, decreasing the net yield of sheep in total. Everyone would be better off if they cooperated.

There are two potential approaches to preserving the common space. One is to divide it up and assign property rights to individuals. In this case, each individual would have the incentive to preserve his piece of the field by not overgrazing because that individual would directly benefit from careful stewardship.

However, for land stewardship, private ownership is not a universal environmental panacea because some element of commonality remains. The aquifer holding water underneath the ground may be over-tapped, or soils may be depleted throughout the lifetime of the individual (a sort of commons over time). For fisheries, the ownership society system has holes in it because fish swim around. For the case of global warming, dividing the atmosphere up into parcels is impossible by the nature of the problem. The other potential solution is for some form of collective self-regulation (see Negotiations section).

The decisions made by a single individual are often irrational and therefore not really well described by economic theory. The collective decisions made by whole societies of individuals, in contrast, often exhibit great sagacity. Economics aficionados speak with reverence of the market, economist Adam Smith's "invisible hand." The market, with its law of supply and demand, is a negative feedback system, just like the many natural examples in Chapter 7.

However, the market has a blind spot, which produces the tragedy of the commons. In economics, a cost that is not paid by the decision maker is called an **external cost**. An example of an external cost associated with driving to work is the traffic. One more car will tend to slow down all the other cars on the road, costing other drivers their time. The cost of climate change is not only paid by people who are responsible but by everybody, soon and far into the future. Our sheep farmer made the other farmers pay part of the cost of his new sheep. If true costs are left external, then the economic balancing act of the market does not take them into account, and tragedies of commons are the result. An external cost can be internalized by means of taxes or regulations. The idea is to make the market aware of the true cost of a decision to take this path versus that path.

Cap-and-trade and carbon tax would both tend to bring the cost of common environmental damage into a private decision to emit CO₂.

One way to harness the balancing abilities of the market for preventing global warming is a scheme called **cap-and-trade**. A regulatory agency allocates permits for emission of climate forcing agents such as CO₂, with the total number of permits they issue totaling some lower overall rate of emission than business as usual. If a company or a country is able to cut its emissions to even lower than its allocation, it has the right to sell its allocation in what are called **carbon offsets**. Some companies already exist whose function is to sell carbon offsets,

using the money to find ways of cutting CO₂ emissions. The price of the permits is set by supply and demand, the same as any other commodity. The industry for which it is most expensive to reduce emissions would be willing to pay the most for the emission permits. It will induce other industries to reduce their emissions because they can make money doing it, selling their permits to the highest bidder.

In this way, the market finds the most efficient, that is to say, cheapest, means of limiting global CO₂ emissions to the cap value. A cap-and-trade scheme for limiting power-plant sulfur emissions (to cut down on acid rain) in the United States has worked well, cutting emissions far more cheaply than anyone had projected. A complication with cap-and-trade for acid rain is that it may allow all the pollution to concentrate in one region, which may be too much for that region. CO₂ is a truly global problem, in contrast, and it makes no difference where the CO₂ is released, so CO₂ emissions would be an ideal problem for limitation by cap-and-trade.

A disadvantage of cap-and-trade is that it seems inherently “slippery,” as if it could reassure people without actually solving the climate change problem. One can buy carbon offsets today on various private markets, and they are absurdly cheap. An American can theoretically eliminate his or her carbon footprint by buying a few hundred dollars worth of carbon offsets per year. However, if the entire country wanted to go carbon neutral, it would cost much more than a few hundred dollars per person. The price discrepancy is because the carbon offsets available now are based on what is called “low-hanging fruit,” the cheapest carbon emissions to eliminate. So, in an overall sense, the buyer of the offsets is not really paying the true cost of their carbon emission. Cap-and-trade might also encourage richer nations to continue emissions because they can afford to do so.

The alternative strategy is to tax carbon emission, using the tax code to internalize the costs of the environmental damage. In the end, carbon tax and cap-and-trade accomplish the same thing, but the tax specifies the price and allows the market to choose the emission, whereas cap-and-trade sets emissions and lets the market choose the price. A carbon tax is easier for a company to plan for financially and may have more stability than the effects of cap-and-trade, encouraging the company to invest in reducing emissions rather than speculating on the future of the carbon market. On the other hand, it is impossible to know in advance what tax rate would be required to reach a particular rate of carbon emissions. A carbon tax would have to be tuned repeatedly if it is to achieve a predefined cutback of carbon emissions.

The market may influence the global warming debate through the insurance industry. The last few decades have seen a huge increase in insurance payout to cover atmospheric-related natural disasters. It is not at all clear to what extent this trend is due to climate or to social factors, but the insurance industry is certainly concerned about the possibility of climate change affecting their bottom line. Insurance companies have been responsible for other preventative measures, such as requirements for hard hats on construction sites.

Pollution problems get more difficult to solve, politically, as their footprints expand from regional to global because of the commons effect. Local issues like contaminated drinking water or urban smog are clearly in the interests of local individuals to solve. You fix it, just as you would fix the roof over your head. A clear link exists between the costs of cleaning up and the benefits. Larger-scale, regional problems tend to run into us-versus-them issues. Why should the state of Ohio sacrifice their right to burn high-sulfur coal when the costs of that consumption are paid by the folks in New York? Why should power plants in Chicago pay to clean up their mercury emissions when most of the mercury emission to the atmosphere comes from China? At the most difficult, worst end of this problem, global warming is the most global of issues, the most complicated type to solve.

Global warming is also a difficult issue to address because it is long-term. Values need to be related not only from one place to another, but across time. The way economists deal




Figure 13-2 The discount rate and future costs. How much it would be worth today to avoid a cost of \$100 at some point in the future, assuming a 3% discount rate.

with values across time is via a construct known as a **discount rate**. The idea is based on a question: which would an individual prefer, a cost today or a cost in the future? A cost in the future, certainly, would be the natural response. A rational reason to justify that choice is that one could invest money today and use the interest to help pay the costs in the future. If we were to assume an interest rate of say 3% per year (above the rate of inflation), we could pay a cost of \$100 in 100 years by investing \$5 today (Figure 13-2). It would be worth paying \$4 today to avoid a \$100 cost in 100 years, but it would not be worth paying \$6 today. The bottom line of this idea is that costs tend to shrink as they get further off into the future.

Of course, the idea of anybody actually going out and opening a bank account with \$5 in it, to save for that cost in 100 years, seems rather quaint. However, economic theory is more than just a prescription for how one could rationally make the most of one's assets; it is also a description of how money tends to flow. The discount rate theory illustrates why financial decisions made by the market tend to be somewhat short-sighted. Water flows downhill. Money flows short-term.

Negotiations

In 1988 agencies of the United Nations formed the Intergovernmental Panel on Climate Change (IPCC). The role of the IPCC is to assess “the scientific, technical and socio-economic information relevant to understanding the scientific basis of risk of human-induced climate change, its potential impacts and options for adaptation and mitigation.”

Climate science is such a vast field that it is difficult for a single person to know it all. But it has been summarized and integrated by the IPCC.

The purpose of the IPCC is to publish reports summarizing the scientific literature. These reports are divided into three topic areas, written by different working groups. Working Group I deals with the scientific basis for the climate change forecast. WG II deals with the impacts of climate change on the natural and human world. WG III assesses options for limiting greenhouse

gas emissions or otherwise avoiding climate change. The most recent reports from all three are available on line at <http://www.ipcc.ch/>.

The reports consist of chapters and various levels of summaries. Much of the work of writing and reviewing the documents is done by researchers at universities, not necessarily employed by any government. The rules are that the chapters may present no previously unpublished research but are based only on published, peer-reviewed scientific papers.

Of course, not every paper published in the scientific literature is guaranteed to be correct in every way, but the practice of sending manuscripts out to rival scientists for review makes sure that most of them are at least plausible. If a scientist disagrees with an article that has already been published, he or she is free to write another article attempting to disprove the first one. Articles are sometimes retracted by their authors if they are proven wrong, or sometimes a disagreement just “festers” in the literature, requiring judgment on the part of the readership to decide which is more likely to be right or wrong.

Once an IPCC chapter is drafted, the chapter is subject to another round of reviews by scientific experts and government workers. The chapters are grouped together into books, which also come with technical summaries and summaries for policymakers. These summaries are subjected to line-by-line review, voting, and approval by the entire IPCC Panel. This is a long and painful process!

The conclusion of the first, 1990, IPCC report was that it was not yet possible to detect human-induced global warming. This was largely because the climate models were driven at that time by the observed rise in CO₂ and other greenhouse gases, but neglecting aerosols, the other human impact on climate, and changes in solar intensity; so the model temperatures did not fit reality very well. By 1995 these effects had been added, and the models fit better. In addition, the spatial patterns, or **fingerprinting**, of the temperature changes, such as an observed cooling of the stratosphere, fit the signature expected from greenhouse gas forcing of climate. The 1995 report issued a now-famous statement that “the balance of evidence suggests a discernable human influence on global climate,” thus providing the impetus for drafting the **Kyoto Protocol** in 1997.

Negotiations to limit CO₂ emission began with a document called the **Framework Convention on Climate Change**, or **FCCC**, that was drafted during the Earth Summit meeting in Rio de Janeiro in 1992. The objective of the FCCC was to achieve “stabilization of greenhouse gas concentrations in the atmosphere at a level that would prevent dangerous anthropogenic interference with the climate system.” The FCCC deferred defining what the dangerous concentration would be, a discussion that continues to this day.

Countries have started talking to each other about avoiding climate change, but they haven't done much about it yet.

The FCCC established a procedure for negotiations by setting up a series of meetings between countries, called **Conference of Parties**, or **COP** meetings. The FCCC also established a procedure for creating amendments to itself, the most famous of which was drafted during the COP-3 meeting in 1997 in Kyoto, Japan, called the **Kyoto Protocol**. The Kyoto Protocol attempted to limit global CO₂ emissions to about 6% below 1990 emission levels by the year 2010.

The treaty came into force when countries accounting for 55% of CO₂ emissions globally agreed to its terms. The United States withdrew from negotiations under the Bush administration in 2001, presenting a fairly sizable obstacle to achieving this target, seeing as how the United States accounts for 25% of CO₂ emissions. Even without the U.S., the agreement came into force when Russia agreed to the Kyoto Protocol in 2005. In practice, the agreement has no

real “teeth,” and the rate of CO₂ emission has grown with time rather than declined. In fact, emissions are higher today than the most aggressive growth business-as-usual scenario from the year 2000, called the A1F1 scenario.

The Kyoto Protocol expires in the year 2012, and there are COP meetings every half year or so with the aim of negotiating a successor agreement. So far, progress has been slow. In fact, the **Montreal Protocol**, an agreement to limit Freon emissions to save the ozone in the stratosphere, arguably has done more to slow global warming than the Kyoto Protocol has!

One line of dispute can be drawn between the rich industrialized and developing nations. Historically, most of the fossil fuel emissions have come from industrialized countries, although Brazil comes up as a large historical emitter when carbon release from land-use changes are also considered. The developing nations argue that it would be unfair to base national emissions allowances in the future on past emissions because it would be carving into law the imbalance between industrialized and developing nations. We have seen in Chapter 9 the intimate relationship between CO₂ emission, energy, and economic activity. Within our current technological stage of development, CO₂ emission is the key to national wealth and comfort. The developing world understandably wants these benefits also and can claim that capping their emissions according to current emissions would be unfair.

On the other hand, emissions are growing rapidly in China and India, and China has recently surpassed the United States as the largest national emitter of carbon. The emission rates per person are still much lower in China and India than in the United States. Even so, stabilizing the atmospheric CO₂ concentration will require emission restrictions from the developing world, too; restrictions on the rich nations alone would not be enough.

Economics

One way to approach the decision of whether to limit CO₂ emissions is based on economics. Which path costs less—to prevent climate change aggressively or instead to endure it?

A caveat: the economy is much harder to forecast than climate. Physical sciences are based on a solid foundation of natural laws such as Newton’s equations of motion, conservation of energy, and so on. It is not always possible to derive the physics of climate solely from these natural laws. Higher-level processes like turbulence and clouds are called “emergent behavior,” and are surprisingly difficult to build from the fundamental laws of physics. But at least the foundation exists.

In contrast, economics and the other social sciences are not as securely founded. Economic trends are strongly impacted by technological progress, which is impossible to forecast, and by social and political fancies. One could even argue that economics is pathologically unsuited for forecasting because it is self-aware. If I came up with a model that could successfully forecast stock prices, I could use it to make money—until everyone else discovered my methods. At that point, the stock prices would reflect everyone’s knowledge of the model forecast, and the future trajectories of the stock prices would alter to evade my forecast. I would not make money any longer. Not even on my darkest days do I imagine that the physical world that I model for a living is going to change its behavior for the purpose of making my model wrong!

Economics is an awkward tool for the climate change decision because many aspects of the natural world are simply not economically fungible. The economic value of coral reefs could be tabulated as the tourist dollars they generate, same as Euro-Disney, or as some willingness to pay to save them, today, which might be a very different number from how much the last living reefs would seem to be worth to people in the future. But ultimately, if coral reefs cease to exist because it gets too hot, we cannot purchase more coral reefs for any amount of money.

The bulk economic framework also hides an issue of fairness, in that people who benefit from business as usual are not the same people who pay for business as usual, say by living in an increasingly drought-prone world far in the future. I imagine that giving up the institution of slavery in the United States entailed a cost, too, for somebody. Ultimately the abolition of slavery was not an economic decision, but an ethical one.

In general, the costs of climate change are predicted to be nonlinear with temperature increase. A small increase in temperature, say 1°C or less, will probably not do too much damage and might even be beneficial on the global average. As the temperature rises to 2° or 3°C or higher, however, the economic impact starts to rise sharply. This is true for many aspects of climate change: farming, sea-level rise, health impacts, and the possibility of abrupt climate change.

Although economics is not really well-suited to computing a cost of climate change, it can be used, at least in the short term, to estimate the costs of reducing carbon emission from human activity. The reports written by Working Group III of the IPCC are devoted to this question: the technology and economics of reducing, or **mitigating**, the human impact on climate.

More and cheaper CO₂ emission cuts are possible in the next few decades from improvements in energy efficiency than from new alternative energy sources.

The results of the 1300-page Working Group III report are summarized in Table 13.1, a listing of options for reducing carbon emission from different sectors of the economy. In each category there is a quantity of carbon emission cuts, or of the equivalent amount of some other greenhouse gas, usually methane, available in the next few decades for a specified cost of \$366/ton of carbon (which comes to \$100/ton of CO₂ molecules).

Table 13.1 Potential Cuts in CO₂-Equivalent Emissions between Now and the Year 2030*

Sector	Carbon-Equivalent Emissions Savings, Gton CO ₂ /Year	
Energy supply	1.0	Shifting to carbon-free renewable energy and capturing CO ₂ from coal
Transportation	0.8	Major improvements in efficiency are possible at low cost
Buildings	1.1	Large savings are possible and would pay for themselves
Industry	1.1	This slow-growing sector could be more efficient
Agriculture	0.9	No-till agriculture could capture carbon from the atmosphere
Forestry	0.7	Reforestation
Waste treatment	0.3	Smaller emissions savings, mostly methane and N ₂ O
Total	5.9	

* Available at a cost of \$100/ton of CO₂-equivalent. The other greenhouse gases are included in this total, converted into their equivalent CO₂ emissions based on their relative potency as greenhouse gases. The numbers are from Table TS.17 in the Intergovernmental Panel on Climate Change Working Group III report. Where there was a range of uncertainty in the original table, it has been averaged to the middle value.

Most of the potential emission savings in this time frame come from improvements in energy efficiency rather than from changes in where energy comes from. There are six entries in Table 13.1, all comparable in magnitude of about 1 Gton C per year. One of these is driven by changes in energy supply, and the other five all have to do with efficiency. One of the entries, energy use in buildings (heating, lighting, and air conditioning), actually has negative costs for a substantial fraction of the emissions cuts, changes like improved insulation and lighting, which would save money over the long run rather than costing extra.

In aggregate, averting climate change is projected to cost a few percent of the global domestic production, **GDP**, a measure of economic activity. A few trillion dollars per year is a lot of money, but this cost would probably not be noticeable to people in their everyday lives. It would be comparable to the cost of a major war, such as the American involvement in Iraq. Perhaps the average U.S. citizen would be financially better off if that war had never started, but it is purely hypothetical because that other world does not exist.

The GDP of the Earth is growing every year by an amount comparable to the projected costs of clean energy. A GDP growth rate of 3% per year sounds small, a few pennies out of a dollar, but yet somehow that economic growth has a huge impact on people's livelihood and sense of prosperity. If the economy grows by 5% each year, everyone gets rich and shoeshine boys trade stock tips. The Great Depression decreased global productivity by about 25%. The reunification of Germany put a drain of about 5% on the GDP of West Germany, which had a noticeable impact on life for West Germans.

Let us compare two hypothetical worlds, with and without cuts in CO₂ emissions, in Figure 13-3. The CO₂-abatement cost is like a tax, applied to all the money that we make, every year. When we put together the 3% cost and the 3%/year growth rate in GDP, we see that the effect of the cost would be to set everything back about a year. A person would get only \$97 this year instead of \$100, for example, because of that 3% cost. But by next year, the economy will have grown to \$103, and she would get her \$100, just one year late.

Costs of avoiding climate change would be comparable to the costs of deciding to go to war—significant but not prohibitive.

A 3% cost slows only the actual growth rate of the GDP from our baseline growth of 3%/year to an adjusted growth rate of 2.91%/year. This does not seem like a catastrophe to a naïve noneconomist like this author. To an Earth scientist, any positive exponential growth like this is an explosion, and whether it explodes quickly or slowly, it cannot persist forever.

Although current negotiations are far from the ultimate goal of preventing climate change, a strong agreement does provide the benefit of guiding investment toward decreasing CO₂ emission. Coal power plants, for example, can be built using newer technology called **integrated gasification combined cycle**, or **IGCC**, enabling future CO₂ capture for sequestration into the geological repositories. If CO₂ is released to the atmosphere, IGCC plants cost about 20% more than traditional plants. If CO₂ sequestration starts to be common practice, however, IGCC plants would be cheaper overall than traditional plants. In the regulated U.S. energy market, regulators often require power companies to use the cheapest technologies. In general, this rule makes sense if we want to limit the potential of monopoly power companies to spend money wantonly, passing the costs on to their captive customers. But in this case, the cost of climate change is external to the accounting, and so the regulation may not be choosing the overall cheapest option. Once a power plant is built, it has a life span of 50 years or so. The ultimate costs of decreasing CO₂ emissions would be lowered by starting now rather than continuing to invest in high-emission technology.




Figure 13-3 The effect of an ongoing expense equivalent to 3% of gross domestic product every year on an economy that is growing at a rate of 3% per year.

Dangerous Anthropogenic Interference

Our discussion so far this chapter has centered on what can be done to limit CO₂ emissions at what cost. The other end of the question is, how much climate change is too much? The goal of the FCCC is to “prevent dangerous anthropogenic interference with the climate system,” but the FCCC did not specify what they meant by dangerous.

The most basic metric of the changing climate is the change in the average temperature of the Earth. Negotiators have been using a temperature change of 2°C as a benchmark target. Mostly this temperature change was decided upon by comparison with global temperature changes in the past. It is probably as warm today or warmer than it has been throughout the history of civilized humanity and agriculture. Global mean temperature may have been 1.5°C warmer during the last interglacial period, 120,000 years ago, or the last time the Earth’s orbit resembled the way it is today, a time called Stage 11, which was about 400,000 years ago. A human-induced warming of 2°C would make the Earth warmer than it has been in millions of years. The 2°C limit is not really a “danger limit” in that a temperature change of 1.9°C still seems plenty dangerous. But the projected damages tend to go up rather drastically with warming greater than 2°C, so that amount of warming provides a benchmark to talk about, what it would take to avoid surpassing.

Jim Hansen [2004] argues for 2 m of sea-level rise as a danger limit. The IPCC forecast for sea-level rise is about half a meter in the next century, detrimental but not yet a show-stopper according to the judgment of Hansen. But if the oceans were to rise by 2 m, a serious amount of land would be lost. Paleoclimate reconstructions (Figure 12-8) seem to predict that in thousands of years, sea level could rise by 10s of meters from doubling of CO₂. Hansen argues that the paleoclimatic record would suggest the possibility that 1°C temperature increase beyond today might exceed his 2-m sea-level danger limit.

Another issue is how fast the climate changes. The FCCC document states that warming should be slow enough for ecosystems to adapt naturally, with no threat to food production, “and to enable economic development to proceed in a sustainable manner.” In other words, the transition to a new climate regime should be slow enough for natural and human systems to adapt. Here again, history can be a guide. The trick is that temperature changes in the past tend to be faster for short periods, but when they are averaged over longer periods, the rates of change tend to be slower. Comparisons must be made over similar stretches of time.

The current rate of warming globally over the past decade is 0.2°C per decade, higher in some locations. The rate of change averaged over multiple decades drops to 0.1°C per decade on average. Historically, warming at 0.1°C per decade is a fairly unusual warming. Based on this, a science advisory ground, the German Advisory Council on Climate Change (<http://www.wbgu.de>) concludes that 0.2°C per decade seems about the upper safe limit to the long-term temperature rise.

The next tasks are to figure out what the climate change targets mean in terms of atmospheric CO_2 concentration, and how much CO_2 can be released to stay within that limit. As a first step, the climate sensitivity can be used to predict the eventual warming that would result from increasing atmospheric CO_2 for a long time, say a few centuries (Chapter 4). According to Equation 4-1, a temperature change of 1°C is already in the works from our pCO_2 concentration of 385 ppm; 2°C would be generated by pCO_2 of about 450 ppm, and 560 ppm would generate about 3°C (Figure 4-7).

To relate atmospheric concentrations to the allowable emissions, the initial approach was to construct what are called **stabilization scenarios**, somewhat arbitrary but reasonable-looking trajectories for the CO_2 rise in the atmosphere, ending up at stable concentrations of 350, 450, 550, 650, or 750 ppm.

A few decades ago, scientists seemed to think of 550 ppm as a good CO_2 target, about double the natural value of 280 ppm, but this has been ratcheting down as the climate has actually started to change. Limiting the warming to 2°C would require a limit of roughly 450 ppm. Now many scientists are embracing 350 ppm as a real “danger limit” CO_2 concentration, as articulated again by Jim Hansen. The idea is that the Earth would continue to warm even if CO_2 remained at its current 385 ppm and did not rise any further, as the so-called committed warming materializes. Dangerous changes in climate are already under way, his reasoning goes, in the form of droughts and general storminess. Preventing the committed further warming would require backing the atmosphere down from 385 to 350 ppm.

We may already be past a reasonably defined “danger limit.”

Carbon cycle models can be used to predict how much carbon would be taken up into the ocean and the land surface. But it was discovered that the details of the carbon emissions you could get away with, whether cuts have to be made this year, versus in the next decade or several decades into the future, depends a lot on the details of the prescribed CO_2 rise in the stabilization scenario. It was very confusing; the scenarios did not really answer the more pressing question of how deeply carbon emissions had to be cut, and how soon, to avoid exceeding the climate change benchmark.

The picture gets clearer if we frame the question in terms of the total human-released carbon slug over the entire fossil fuel era and just ask about the maximum warming at the peak. A lot of the complexity of the carbon cycle and climate system cancels out this way.

One piece of the story is that if CO_2 is released more quickly, the oceans and land surface would lag behind in absorbing it, so the airborne fraction of the released CO_2 would be higher. A higher airborne fraction tends to increase the warming you get.

Compensating for this is the time lag between rising CO_2 and the changing climate. It takes decades, even centuries, for the climate of the Earth to respond fully to a change in atmospheric CO_2 concentration, mostly because of the thermal inertial of the oceans. A fast CO_2 release would have a somewhat lower climate impact than a slower one because the CO_2 reaches the peak and starts back down sooner for the fast release.

When these two competing effects are put together, it turns out that they more or less cancel, so that the peak temperature change that the Earth will ever see does not really depend on how quickly the carbon is released. Framing the question this way does not answer the question of how much carbon needs to be cut today versus in a decade or a few decades. But

it says that if carbon is emitted now, it cannot be emitted later, a zero-sum game that could be used as the basis for carbon emission regulation.

Slug theory—the amount of warming at the peak warming is determined by the total amount of CO₂ ever emitted by humankind.

So far, humankind has combusted about 300 Gton C worth of fossil fuels and deforested about another 200 Gton C worth of trees, for a total of about 500 Gton C. It turns out that limiting the peak warming to 2°C requires limiting the total slug of carbon to about 1,000 Gton C, more or less. So we are halfway there.

Our current rate of carbon emission is about 10 Gton C/year, so we could go for perhaps 50 years if the rate of carbon growth held steady, and then it would have to drop immediately. But the rate of CO₂ emission is growing rapidly, projected to double by midcentury under business as usual. In that case we would run into the limit sooner, say in 30 years. This does not mean that there are 30 years before we need to worry, because it would be hard to stop emissions “cold turkey” at that time.

If instead we imagine that at some point in the future the rate of CO₂ emission begins to decline at some rate, the trajectories would look like Figure 13-4a. The total slug of emitted




Figure 13-4 Rates of CO₂ emission cuts required to keep the total slug of CO₂ ever released (the area under the curve) under 1,000 Gton C. The *bottom panel* shows the required rate of decrease plotted as a function of the year in which the cuts begin.

CO₂ is proportional to the area under the emission curve. Start now, and we could get there with cuts of 2% per year, but if growth continues for 20 years, the cuts have to be sharper (Figure 13-4b). Because cutting emissions requires changing infrastructure, faster cuts would cost more. The situation is like a term paper assignment for a class that has some flexibility in the due date, but if you hand it in late it has to be a longer paper.

Alternatives

The online model for the Kaya identity in Chapter 9 (<http://understandingtheforecast.org/Projects/kaya.html>) can be used to estimate how much carbon-based energy would have to be replaced by carbon-free energy to achieve atmospheric stabilization.

The Kaya model makes a forecast of energy use and carbon emissions based on past trends, what is called business as usual. The resulting CO₂ rise is shown in the middle panel of Figure 9-14, along with atmospheric CO₂ trajectories of the CO₂ stabilization scenarios. The bottom panel of the figure shows how much carbon-free energy would be required to get from business as usual to each of the stabilization scenarios. The calculation assumes that the carbon-free energy is replacing coal. To stabilize at 450 ppm would require about 17 terawatts (TW) of carbon-free energy by the year 2100. For comparison, mankind is currently consuming about 13 TW of energy.

The greatest short-term opportunity is conservation. The United States uses twice the energy per person as is used in Europe or Japan (Figure 9-11). In part, this is because the human terrain of the United States developed in a more automotive way, whereas public transit and trains are better developed in Europe and Japan. In addition, Americans drive larger cars and live in larger houses.

Most of the cheapest solutions involve energy efficiency, rather than
new energy sources.

On the other hand, it seems unlikely that conservation alone will solve the world's energy problems. Most of the world lives at a standard far below that of Europe or the United States. If the citizens of the developing world ultimately consume as much as the American model, the footprint of humanity on the Earth would increase by about a factor of 5. Projections of energy use in the future already assume continued increases in energy and carbon efficiency, and still predict a thirst for energy that results in higher CO₂ concentrations than the stabilization scenarios.

None of our traditional sources of carbon-free energy seems capable of making up the shortfall for the coming century by itself (Hoffert et al., 2004; Pacala and Socolow, 2004). So the strategy has been devised of dividing up the emissions cuts into manageable pieces called “**wedges**,” each of which ramps up from nothing today to savings of something like 1 Gton C per year by 2050 (Figure 13-5).

The first option that seems to come to people's minds is nuclear energy because it generates no carbon, and we know it works. The French are using nuclear for more than a third of their electricity needs. However, if we want to generate 17 TW of nuclear energy, using present-day standard 1000 MW reactors, it would require 17,000 new reactors within 100 years, for an average reactor construction rate of about one new reactor every second day, continuously, for 100 years. That seems like a lot of nuclear reactors, or else reactors much larger than today's.




Figure 13-5 The difference between business-as-usual emission trajectory and one that stabilizes atmospheric CO₂ has been divided into manageable pieces called “wedges.”

Wedge theory — no single change in infrastructure can accomplish the goal.

It is a close call whether or not there would be enough mineable uranium to support this level of nuclear energy production. The global inventory of mineable uranium today would be sufficient to generate 10 TW of energy at our current efficiency for a few decades before it would be exhausted. However, uranium today is rather inexpensive, and we bother to mine only the purest uranium deposits because production of nuclear energy is limited by supply of reactors, not by supply of uranium. We could move to lower-purity uranium deposits. Ultimately it may be possible to extract uranium from seawater, where it is present in large amounts but at very low concentrations.

In addition, it is theoretically possible to stretch the energy yield of natural uranium by converting it to plutonium in what is known as a **breeder nuclear reactor**. This increases the energy yield from the uranium by about a factor of 50, but the downside is that it is very easy to produce nuclear weapons from plutonium. For this reason, breeder nuclear reactors are not used today.

Windmills are becoming economically competitive with new traditional power plants. Windmills supply 7% of the energy needs of Denmark with no adverse impact on the beauty of the landscape of Denmark in my opinion. Wind energy currently accounts for 0.3% of energy globally and is growing at 30%/year. At this rate, wind could supply 10% of the world’s energy within the next couple of decades. Scaling up current wind energy production by a factor of 50 would generate about 2 TW of energy.

Both surface wind and solar power share a technical difficulty that energy must be stored because periods of maximum power demand may not coincide with times of maximum windiness or sunshine. To some extent, solar and wind technology complement each other in that when the sun is not shining, the wind is more likely to be blowing. Another way to deal with this problem is to improve the power grid so that electricity could be efficiently transported farther around the world.

It may be possible to extract energy from winds at high elevations in the atmosphere. Winds get faster with altitude in the troposphere, peaking in the jet stream in mid-latitudes of both hemispheres. The air up there is at lower pressure than down here, and so it therefore

has less density to drive wind blades, but still the power density is much higher at high altitude than it is at the Earth's surface. Wind energy collectors could function like giant kites tethered to the ground by electrically conducting cable, remaining aloft passively by riding the wind. It has been proposed that high-altitude windmills could scale up to 10s of TW of energy (<http://www.skywindpower.com>).

The sun bathes the Earth in energy at a rate of about 173,000 TW, enough to power all of our energetic ambitions 10,000 times over, but the difficulty is that the power is spread out over the entire Earth's surface. Photosynthesis manages to capture about 100 TW globally.

There are two technologies for extracting electricity from sunlight, solar thermal and photovoltaic. **Solar thermal** is the lower-tech of the two, concentrating the solar power with mirrors into a stream of water, which is converted to steam and used to power an electrical generator. It is similar to a traditional coal plant except that it uses the sun instead of coal to make steam. Because it relies on mostly established industrial technology, commercial-scale solar thermal plants can begin construction now. One downside of both wind and solar power is that they are intermittent, dependent on the weather. However, the heat from a solar thermal plant could be stored in a reservoir such as molten salt. Heat could be extracted from the reservoir to generate power even at night or on cloudy days.

Solar cells, or **photovoltaics**, generate electricity directly from sunlight using semiconductor technology. Photovoltaics are currently rather expensive to produce, but if solar energy generation were scaled up by a factor of 700 from present-day rates, they could generate several terawatts of energy. Similarly to surface windmills, photovoltaics seem unable to solve the entire energy problem, but they could certainly help.

In the long run, to the year 2100, large new carbon-free energy sources will be required, of a kind that do not exist yet.

One cool idea is to build solar cells on the moon (Hoffert et al., 2002), beaming energy back to the Earth as microwaves. The moon would be an ideal location for solar cells because it never rains there, and there are no clouds and no birds. The cells could be produced from lunar soil, reducing the amount of mass that would have to be lifted out of Earth's gravity by rockets, a huge savings in energy and cost. The energy from the solar cells could be transmitted to Earth via microwaves; this is apparently feasible, safe, and relatively efficient. An array of solar cells in orbit around the Earth would also share the benefits of a clean environment and solar flux that is unimpeded by the atmosphere. Solar cells in orbit, however, have to be produced on Earth and lifted to orbit.

You have probably heard of the **hydrogen** economy. Hydrogen itself is not a primary energy source. There are no mineable sources of energy as hydrogen like the energy in fossil fuel deposits. Hydrogen is produced from other sources of energy. For example, electricity can be used to break water into components—hydrogen and oxygen gas—a process called **hydrolysis**. Hydrogen can also be produced chemically by **gasification** of coal.

Hydrogen could be used to store energy for use in transportation. Because hydrogen is a gas, it must be contained under pressure, as opposed to gasoline, which is liquid at room temperatures, or propane, which can be liquefied at moderate pressure. However, methane is also a gas, and it is already being used to power buses and cars. Hydrogen is more explosive than methane, but in crash tests, this added danger is offset by the tendency of hydrogen to rise up, like helium would, escaping from a crashed vehicle. Hydrogen burns with an invisible ultraviolet and infrared flame unless the flames happen to hit some other material that emits visible light.

Another option for carbon-free power is to use fossil fuels and then dispose of the carbon, a strategy called **CO₂ sequestration**. CO₂ can be captured from the exhaust stream of a power plant or even captured chemically from the free atmosphere. The CO₂ can then be injected and stored in particular geological formations in the interior of the Earth. Methane gas can be found in deep Earth reservoirs that are hundreds of millions of years old, so the idea is, in principle, possible.

Old oil wells are the most immediate possibility, but they have limited capacity and might leak through undocumented old drill holes. The largest type of geological formation that would fit the bill is called **saline aquifers**. These deposits contain water in their pore spaces, but the water has salt dissolved in it, so the assumption is that this water is no good for anybody and we might as well inject CO₂ into it. These aquifers are thought to have the capacity to store 10,000 Gton of C as CO₂, enough to hold all of our fossil carbon. However, some places on Earth, such as Japan, have no access to these types of rock formations.

Scientists are also discussing the possibility of sequestering CO₂ in the deep ocean. CO₂ released into the atmosphere will ultimately mostly dissolve in the ocean anyway; about 75% of the carbon would be in the ocean after the air and water reach equilibrium after hundreds of years, leaving an airborne fraction of about 25%. One could envision direct injection into the ocean as simply bypassing the transit of CO₂ through the atmosphere. CO₂ released in the ocean would equilibrate with the atmosphere toward the same 25% airborne fraction after hundreds of years. Ocean injection would cut the head off the peak, but it would have no impact on the long tail of the extra CO₂ in the air (Chapter 10). The immediate difficulty with ocean injection is that the CO₂ is very concentrated right near the injection site, acidifying the water and killing marine life.

You may have read about an idea called **ocean fertilization**, which has the goal of stimulating the plankton in the ocean to take up CO₂ as a means of carbon sequestration. The idea is superficially attractive in that plankton in many parts of the ocean, in particular in the cold surface waters around Antarctica, are starving for tiny amounts of the element iron. Iron is typically supplied to the surface ocean in dust blown through the atmosphere, but the Southern Ocean is remote enough that not much dust blows down there. When oceanographers add iron to surface water in the Southern Ocean, the plankton bloom. One atom of iron could theoretically allow plankton to take up 100,000 atoms of carbon (more or less), so iron could be a powerful lever in the carbon cycle. The idea is that the plankton would grow, die, then sink to the deep ocean, extracting CO₂ from the surface ocean, ultimately from the atmosphere, and transporting it to the deep ocean for us.

As an ocean chemist, I do not buy this as a way to solve the climate problem. The phytoplankton do bloom in the fertilization experiments, but the extra plankton has not been observed to sink to deeper waters. And even if a large region of the ocean could be fertilized until the plankton are limited by other nutrients, like nitrogen and phosphorus, models of the carbon cycle do not predict much drawdown in the CO₂ concentration of the atmosphere. The reason is that it takes hundreds of years to reach CO₂ equilibrium between the atmosphere and the ocean. Ocean fertilization could have a larger impact on atmospheric CO₂ if we were willing to wait 500 years, but the impact in a few decades is too small to be worth it. Also, it would be impossible to say how much CO₂ was removed from the atmosphere, making the technique too opaque to use in a financial carbon offset market.

The last idea to mention is to cool the climate of the Earth deliberately by **geoengineering** to counteract the warming from CO₂. One possibility is to generate a haze in the stratosphere by releasing sulfur or some other chemical. Aerosols cool the Earth by scattering sunlight, as the sulfate aerosols in the troposphere from coal combustion do today (Chapter 11). The advantage of putting particles in the stratosphere is that no rain originates from there, so the lifetime of the injected aerosols would be several years, as opposed to several weeks in the troposphere.

The volcanic eruption of Mt. Pinatubo in 1991 injected a haze into the stratosphere and cooled the planet down by about 0.5°C for more than a year. The aerosols could be shot up to the upper atmosphere in shells fired by large cannons, or sulfur could be added to the jet fuel of commercial airliners.

Another proposal for climate engineering includes the placement of a large reflective object in space, in orbit around the Earth or at the stable “Lagrange” point between the Earth and the sun where objects can set indefinitely. The space-based option would be more costly and would require active removal if its effects were eventually judged to be detrimental, and it would require ongoing management and maintenance to continue working.

My objection to these proposals is that the warming effects from CO_2 will last for millennia, committing humankind to taking an active role in climate maintenance essentially forever in exchange for a few decades of careless energy production. If the climate management in the future broke down for whatever reason, the brunt of the higher CO_2 concentrations would kick in very quickly. We would be leaving the Earth on life support. I could imagine emergency geoengineering efforts to cool the climate down for a few decades, buying time for natural or human-driven CO_2 extraction from the atmosphere. But geoengineering cannot be considered a “fix” for the climate problem: it could only be a temporary Band-Aid.

TAKE-HOME POINTS

- Human-induced climate change is an example of the tragedy of the commons. Economists refer to climate change as an external cost, which can cause the market to make poor decisions. Economic forces also keep the market focused on short-term profit and cost, disregarding costs that come far in the future.
- International negotiations under the auspices of the United Nations have not made much progress so far in actually reducing greenhouse gas emissions.
- The ultimate goal of negotiations is to prevent “dangerous interference” with the climate system. Danger, in this context, might be benchmarked in terms of global temperature, rate of warming, sea-level rise, peak atmospheric CO_2 concentration, or the total amount of CO_2 ever released.
- The sooner we begin changing the energy infrastructure, the cheaper it will be to avoid dangerous climate change.

STUDY QUESTIONS

1. What is an external cost in economics? How does it relate to the tragedy of the commons?
2. Why is the discount rate an important factor in people’s view of the economics of avoiding climate change?
3. How can danger limits for climate change be defined? What specific implications do they have for carbon emissions in the future?
4. Which can provide more carbon emission reduction in the next few decades, efficiency or alternative energy sources? What about to the year 2100?

FURTHER READING

- German Advisory Council on Global Change (WGBU). *Climate Protection Strategies for the 21st Century: Kyoto and Beyond*. Special Report, 2003.
- Hansen, J. Defusing the global warming time bomb. *Scientific American*, 290:68–77, 2004.
- Hoffert, M. I., K. Caldeira, G. Benford, et al. Advanced technology paths to global climate stability: energy for a greenhouse planet. *Science*, 298:981–987, 2002.
- Hoffert, M. I., K. Caldeira, A. K. Jain, E. F. Haites. Energy implications of future stabilization of

- atmospheric CO₂ content. *Nature*, 395:881–884, 1998.
- Nordhaus, W., J. Boyer. *Warming the World: Economic Modeling of Global Warming*. Cambridge, Mass.: MIT Press, 2000.
- Nordhaus, W. Global warming economics. *Science*, 294:1283–1284.
- Pacala, S., R. Socolow. Stabilization wedges: solving the climate problem for the next 50 years with current technologies, *Science*, 305:968–972, 2004.

Exercises

1. **Compound Interest.** The formula to compute compound interest for a bank account is

$$\text{Balance}(t) = \text{Balance}(\text{initial}) \cdot e^{k \cdot t}$$

This was first presented in Chapter 5, when we met the exponential function e^x . Assuming an interest rate of 3%/year, how much would an investment of \$1.00 made today be worth in the year 2100? What if the interest rate were 5%/year?

2. ΔT_{2x} . The formula to estimate temperature response to changing CO₂ concentration is

$$\Delta T = \Delta T_{2x} \times \frac{\ln(\text{new } p\text{CO}_2 / \text{orig. } p\text{CO}_2)}{\ln(2)}$$

This formula was first introduced in Chapter 4. Calculate the temperature that would result from increasing atmospheric pCO₂ from pre-anthropogenic value of 280 ppm to 1,000 and 2,000 ppm. The temperature during the Cretaceous period might have been 6°C warmer than today. Using a ΔT_{2x} value of 3°C, what must the pCO₂ have been at that time? How does your answer change if ΔT_{2x} is 4°C?

3. **Carbon-Free Energy.** The Kaya Identity Web page actually runs a carbon cycle model to predict the atmospheric pCO₂ response to its predicted carbon emissions. You learned about this model in Chapter 9. The Kaya Web page then computes how much coal energy would have to be replaced by carbon-free energy, if we wish to stabilize atmospheric CO₂ at some concentration (choices are 350, 450, 550, 650, and 750 ppm). Using the default Web page settings, which are something like a business-as-usual scenario, find from the plot the amount of energy in terawatts required to stabilize CO₂ at 450 ppm.

- a) If a typical nuclear reactor generates 1000 megawatts of energy, how many power plants would be required by the year 2100? (The prefix “tera” means 10¹², and mega means 10⁶). How many power plants would this require?
- b) A modern windmill generates about 1 megawatt of energy; let us say that future ones will generate 10 megawatts per tower. How many of these would be required to meet our energy needs by 2100? The radius of the Earth is 6.4 × 10⁶ m. What is its surface area? Land occupies about 30% of the surface of the Earth; what area of land is there? Let us assume that windmills could be placed at a density of four windmills per square kilometer. What fraction of the Earth’s land surface would be required to supply this much wind energy?

4. **Carbon Stabilization and Kyoto.** How much would we have to cut emissions to stop atmospheric pCO₂ from rising beyond current levels? You could just randomly plug numbers into the model, or you could do it in a smarter way: Run the model for BAU (climate “geekspeak” for business-as-usual) and determine from the printed output below the plots what the rate of CO₂ uptake is by the ocean today (call it year 2000). You will have to take the difference in the cumulative ocean inventory between two adjacent time points and divide by the number of

years between those two points (five), to get Gtons of C uptake per year. Plug these numbers into the model and run it to see whether that really stops $p\text{CO}_2$ from going up. What percentage change in emission is this? The Kyoto Protocol aimed to limit emissions to some percentage below 1990 levels. The Japanese and Europeans argued, at the high end, for something like 9% reductions below 1990. The Americans and the Saudis were at the low end, say 2% below 1990. Compare this range of emission with your answer.

Glossary

Absolute zero—The lowest possible temperature when atoms have the minimum possible amount of vibrational and translational energy. This is 0 K on the Kelvin scale, -273.15°C on the centigrade scale.

Acid—A water solution containing high concentrations of H^{+} ions.

Acid rain—The effect of sulfate aerosols, consisting of sulfuric acid, on the acidity of rainwater.

Active zone of permafrost—Surface soil that melts in summer, with ice below.

Adiabatic—Insulated from the surroundings, unable to gain or lose heat from the environment.

Albedo—The fraction of incoming solar radiation that is reflected back to space without ever being absorbed. Ice and clouds have a high albedo, so they cool the Earth.

Anaerobic—Without oxygen.

Anthropogenic—Arising because of people.

Atmospheric window—Between 700 and 900 cm^{-1} (wave numbers), there are no greenhouse gases that absorb and emit infrared light. If there are no clouds, light emitted at the ground goes right out to space.

Band saturation—If there is a high enough concentration of a greenhouse gas, light within some frequency range that is absorbed by the gas will be completely absorbed. If the concentration of the gas is increased further, it will not absorb any more light of that frequency. The main absorption band of CO_2 is saturated, so adding one molecule of CO_2 has less impact on climate than does one molecule of methane.

Benthic—Living on the sea floor.

Blackbody radiation—Emission of light from an object at some temperature greater than absolute zero. An object that is capable of absorbing and emitting all wavelengths of light is called a blackbody, and the spectrum of light that is emitted is called a blackbody spectrum.

Borehole records—The temperature in the interior of an ice sheet or in the ground reflects the temperature of the air today and also of the past. Borehole temperatures are one type of paleo-temperature proxy record.

Carbohydrates—Carbon in the intermediate oxidation of building blocks of CH_2O .

Centigrade—Metric temperature scale. 0°C is freezing, and 100°C is the boiling point of water; -273.15°C is absolute zero or 0 K on the Kelvin scale.

Chaos—When small errors or differences grow with time. Because weather is chaotic, it is impossible to forecast weather for more than a few weeks into the future.

Cirrus clouds—Wispy clouds consisting of ice particles in the upper troposphere.

Climate—An average year of the attributes of weather like temperature and rainfall taken over some years.

Climate forcing—Some agent that can alter climate, such as greenhouse gas concentrations, solar intensity, volcanic dust, or anthropogenic aerosols.

Climate sensitivity—The amount of warming, averaged over the whole Earth and the whole year, that would ultimately result from doubling CO_2 . This is abbreviated as $\Delta T_{2\times}$.

Cloud condensation nuclei—Particles in the atmosphere around which water droplets can form. If condensation nuclei are scarce, the cloud will have fewer droplets, each of which is larger. Sulfate aerosols act as cloud condensation nuclei, leading to what is known as the indirect effect of sulfate aerosols.

- CO₂ fertilization—Stimulation of plant growth by rising CO₂ levels. Higher CO₂ levels may ameliorate water stress for plants.
- CO₂ sequestration—The idea of pumping CO₂ back down into the Earth or into the ocean rather than emitting it to the atmosphere.
- Conduction—Transfer of heat through matter by physical contact.
- Contrails—Condensation trails left by aircraft flying in the upper troposphere visible as lines in the sky. After a few hours or days, these may be indistinguishable from cirrus clouds.
- Convection—When a fluid is heated from below or cooled from above until there is dense fluid on top of less-dense fluid, so that it turns over.
- Coral bleaching—Corals under stress may expel the symbiotic algae that they harbor, leading to loss of color. This is often followed by death of the coral.
- Coriolis acceleration—A fake force added to the equation of motion to make the equation work within a rotating frame of reference.
- Cumulus clouds—Tower-type clouds formed by localized upward air motion. These are found in thunderstorms and in the tropics.
- Cyclonic—The direction that air flows around a region of low pressure, such as is found in a storm. In the northern hemisphere, cyclonic flow is clockwise when viewed from above.
- Discount rate—Also called an interest rate. In economic theory, a cost in the future is discounted, with the idea that one could invest a smaller amount of money today and, using the interest, pay off the cost in the future.
- Dissolved inorganic carbon—Abbreviated as DIC, this is CO₂, HCO₃[−], and CO₃^{2−}. There are about 48,000 Gton of C in the ocean DIC pool.
- Eccentricity cycles—Cycles in the circularity or elliptical character of Earth's orbit. Eccentricity varies on timescales of 100,000 and 400,000 years.
- Ecological succession—The process by which an ecosystem changes through time, one group of species followed by another, until a climax ecosystem is reached.
- El Niño—An oscillation of the structure of the atmosphere and ocean, centered in the equatorial Pacific but with climate impacts around the world.
- Electromagnetic radiation—Coupled oscillations of the electric and magnetic fields that together propagate through space. Visible light is electromagnetic radiation, as is infrared, ultraviolet, x-rays, and others.
- Emissivity—Written as ε (the Greek letter epsilon), a constant describing how good a blackbody an object is. If $\varepsilon = 1$, the object is a perfect blackbody.
- Energy intensity—Energy use divided by economic production. The amount of energy it takes to make a buck.
- Ensemble—A collection of climate or weather model runs, beginning from slightly different initial conditions. When an ensemble of runs is averaged, the result is more reliable as a forecast.
- Exponential function—Based on the number e raised to a power x , that is, e^x . The exponential function grows or decays proportionally to the value at any time. Interest for a bank account, population growth, radioactive decay, and light absorption all follow the exponential function.
- External cost—A cost associated with some action that is not paid by the decision maker. This leads to the tragedy of the commons.
- Feedback—A loop of cause and effect. A positive feedback tends to amplify an initial change, whereas a negative feedback stabilizes the system.
- Foraminifera—Single-celled protista in the ocean that secrete shells of CaCO₃ from which paleoclimatologists can produce paleo-temperature proxy records.
- Framework Convention on Climate Change (FCCC)—A procedure for international negotiations to limit CO₂ emissions.
- Frequency—The number of “somethings” per time. The frequency of light is cycles of the electric or magnetic field per time. The frequency of hurricanes is the number per year.

- Geopotential surface—A surface that is flat relative to local gravity field.
- Geostrophic flow—The way that fluid in the atmosphere and ocean flows in equilibrium on the rotating Earth. Geostrophic flow is at right angles to the direction of forcing from pressure or wind friction. In the northern hemisphere, water flows 90 degrees to the right of the wind direction.
- Gigaton—Billion metric tons, equal to 10^{15} g.
- Glaciers—Ice rivers that flow down mountain valleys, originating as snowfall on the mountain.
- Gradient—The amount of change in a quantity, such as pressure or a chemical concentration, per meter of distance.
- Greenhouse gas—A gas that is capable of absorbing and emitting infrared light, for example, CO_2 , water vapor, methane, and ozone.
- Hadley circulation—A planetary scale circulation consisting of rising air at the equator, poleward flow to 30 degrees N and S, subsidence, then equator-ward flow at the surface. The Hadley circulation generates a general pattern of rain at the equator and deserts at 30 degrees N and S.
- Heinrich event—A massive iceberg discharge in the North Atlantic during glacial time, documented by deposition of rocks on the sea floor.
- Humic acid—A form of organic carbon in soils, from the decay of living materials.
- Hydrocarbon—Carbon in the reduced form, chains with hydrogens.
- Ice albedo feedback—Warming leads to melting of ice, lowering of the albedo, leading to further warming.
- Ice sheet—Thick flowing ice that is grounded on land. Greenland and Antarctica are examples today.
- Ice shelf—Floating ice, sometimes hundreds of meters thick, from an ice sheet that flows onto the water.
- Ice stream—A region of the ice sheet where the ice flows much more quickly than the surrounding ice.
- Infrared light—Abbreviated as IR. Longer wavelength light than visible light. IR is emitted by objects at near room temperature.
- Inorganic carbon—Carbon that is fully oxidized. Types of oxidized carbon include CO_2 , CaCO_3 , HCO_3^- (bicarbonate ion), and CO_3^{2-} (carbonate ion) in seawater.
- Intensity—Used in this book to describe energy flow per unit area, W/m^2 .
- Intergovernmental Panel on Climate Change—Abbreviated IPCC.
- Isotopes—Atoms that have the same chemistry but different masses. The composition of stable isotopes can be affected by temperature, so isotopic measurements can be used to create paleo-temperature proxy records.
- Joule—A unit of energy, equal to 4.18 cal, where a cal is the amount of energy it takes to warm 1 g of water by 1°C .
- Juvenile carbon—Carbon that is released from the Earth's interior for the first time in the history of the Earth.
- Kelvin—The absolute temperature scale. 0 K is absolute zero. 273.15 K is the freezing point of water, which is defined as 0°C . One degree Kelvin equals one degree centigrade.
- Kerogen—A form of organic carbon in soils, from the decay of living materials.
- Kyoto Protocol—An international agreement to limit CO_2 emissions.
- Lapse rate—The decrease in temperature of the atmosphere with altitude.
- Last Glacial Maximum—30,000 to 18,000 years ago, abbreviated LGM. Global mean temperature was perhaps 6°C colder during this time.
- Latent heat—Heat carried by water vapor, released when the water vapor condenses to liquid.
- Little Ice Age—1650–1800. A period of generally cool and rather unstable climate in Europe and North America, coincident with the Maunder Minimum.
- Maunder Minimum—1650–1700. A period during which there were no sunspots, coincident with the Little Ice Age.

- Medieval Climate Anomaly—About 800–1200 AD. A period of generally warm, stable climate in Europe, coincident with a prolonged drought in the American Southwest.
- Metamorphic decarbonation—A chemical reaction wherein sedimentary rocks like $\text{CaCO}_3 + \text{SiO}_2$ lose CO_2 to become a silicate rock like CaSiO_3 .
- Methane— CH_4 , the main component of natural gas.
- Moist convection—Convection in the atmosphere with the added complication that water condenses into clouds when the air rises.
- Montreal Protocol—The 1987 international agreement to curtail chlorofluorocarbon emissions.
- Natural gas—Fossil fuel in a gas state, mostly methane.
- Obliquity cycles—The angle of the Earth's axis of rotation relative to the plane of the Earth's orbit.
- Ocean ventilation—Circulation carrying surface water into the deep ocean. Ventilated ocean water has higher concentrations of oxygen. Ventilation carries rising CO_2 into the deep ocean.
- Organic carbon—Carbon that is partially oxidized and partially reduced. Living carbon is generally in this state.
- Oxidized—An atom that donates electrons to another atom in a chemical bond is oxidized. Oxygen is greedy for electrons, so atoms that are bound to it are oxidized.
- Ozone—A molecule consisting of three oxygen atoms. Ozone in the stratosphere protects us from ultraviolet light, whereas ozone near the ground causes respiratory problems.
- Ozone hole—The observed severe depletion of ozone in the Antarctic stratosphere caused by an interaction of chlorofluorocarbons (Freons) and stratospheric nitric acid ice clouds.
- Parameterization—When a climate model is unable to simulate the real physics of some process, such as cloud droplet formation, the process is predicted based on some simplified “informed guess” type of model.
- Permafrost—Soil that has remained frozen for at least the last two years.
- pH—The negative base 10 logarithm of the concentration of H^+ ions in a water solution. A pH value of 7 is neutral, less than 7 is acidic, and greater than 7 basic.
- Photosynthesis—Plants harvesting energy from the sun and storing it as reduced carbon.
- Planktonic—Means free floating as opposed to the swimming “nekton” but in practice, plankton live in the surface ocean.
- Precession cycles—The axis of rotation of the Earth spins relative to the elliptical dimensions of the Earth's orbit. Precession cycles dominate climate variability in low latitudes and take 20,000 years.
- Proxy records—Paleoclimatologists extend their climate records into prehistoric time, using stand-in measurements such as tree-ring thickness, because there were no thermometers back then.
- Radiative convective equilibrium—A state of energy balance for each layer of air in the atmosphere when radiation and convection are both ongoing.
- Radiative equilibrium—A state of energy balance for each layer of air in the atmosphere, when radiation is the only process carrying heat around.
- Reduced—An atom that borrows electrons from another atom in a chemical bond. Hydrogen is loose with electrons, so atoms that are bound to it are reduced.
- Relative humidity—The amount of water vapor in the air, compared with saturation. If relative humidity is greater than 100%, the vapor will tend to condense to liquid, until 100% is reached. If relative humidity is less than 100%, water will tend to evaporate.
- Respiration—Converting reduced carbon to metabolic energy by oxidizing the carbon to CO_2 . Multicellular life uses oxygen for respiration, but bacteria can use sulfate or other compounds in place of oxygen.
- Runaway greenhouse effect—An extreme water vapor feedback that results in boiling dry of the oceans. This happened on Venus, but it is not expected to happen on Earth.

- Scattering—Reflection of light by cloud droplets. This differs from absorption in that the energy of the light is not absorbed and the frequency of the light is unchanged.
- Sea ice—Ice that forms from freezing seawater at the surface of the ocean.
- Silicate rock—An igneous rock. These consume CO₂ when they are weathered.
- Silicate weathering thermostat—Negative feedback that stabilizes atmospheric CO₂ concentration on timescales of half a million years.
- Skin altitude—Average altitude from which infrared energy escapes to space.
- Skin temperature—Temperature, on average, at which the Earth radiates infrared light to space.
- Soil carbon pool—Dead organic carbon, mostly humic acid material. The soil carbon pool contains about 1,500 Gton of C.
- Storm surge—The low pressure in a hurricane or other storm lifts up the sea surface, like a very high tide.
- Stratified—A fluid column with denser fluid on the bottom and less dense fluid on top.
- Stratosphere—A region of atmosphere above the troposphere, from 10 to about 50 km. The stratosphere contains about 10% of the air in the atmosphere.
- Stratus clouds—Layered clouds, formed by broad, gradual upward motion of air.
- Sulfate aerosols—When sulfur is released to the atmosphere from combustion, say of high-sulfur coal, it forms tiny particles of sulfuric acid. These scatter light and act as cloud condensation nuclei.
- Terawatt—10¹² W, abbreviated as TW. Humankind is currently producing and using 13 TW of energy.
- Terrestrial biosphere—Living carbon on land, mostly (by weight) plants. The terrestrial biosphere contains about 500 Gton of C.
- Tragedy of the Commons—A situation arising when the benefits of a common resource accrue to the individual, but the costs of its degradation are shared by all.
- Transient climate response—A standard benchmark for comparing the time-dependent response of climate models to increasing CO₂. The TCR is the temperature when doubled CO₂ is reached, after a ramp-up of 1% per year.
- Tree ring records—Trees grow faster when the temperature is warmer, so the thickness of tree rings is used as a proxy for past temperature variations.
- Tropopause—Boundary between the troposphere and the stratosphere.
- Troposphere—Lower part of the atmosphere, from the ground to about 10 km. The troposphere contains 90% of the air in the atmosphere and is where weather happens.
- Turbulent cascade—Fluid in the atmosphere and ocean flows in eddies or gusts that range in size all the way from the global patterns down to sizes of centimeters or millimeters. Turbulent cascade is the flow of energy from large scales to small, as large eddies break up into smaller ones.
- Ultraviolet light—Abbreviated UV. Higher-frequency light than visible, capable of causing sunburn or worse biological damage.
- Urban heat island effect—Paved cities are warmer than surrounding countryside.
- Viscosity—Resistance against flow caused by friction in a fluid. Molasses is a high-viscosity fluid with lots of internal friction.
- Volcanic CO₂ degassing—CO₂, either from metamorphic decarbonation or juvenile, that escapes from the Earth to the atmosphere.
- Water vapor feedback—An initial warming allows more water to evaporate into the atmosphere. Water vapor is a greenhouse gas, so increasing its concentration leads to further warming.
- Watt—A unit of energy flow, defined as Joules per second.
- Weathering—A chemical reaction wherein a rock, or some of the atoms in a rock, dissolves in water.

Constants and Symbols

c —Speed of light in a vacuum, equal to 3×10^8 m/s.

e —A number used in the exponential function, e^x . The value of e is approximately 2.718.

n —Wave number, defined as the number of waves per centimeter, often used to describe infrared light.

α —Greek letter alpha used here to denote albedo.

ΔT_{2x} — Δ is the Greek letter delta, often used to denote a difference. In this case, it is the difference between two temperatures, an initial and a final after doubling CO_2 .

λ —Greek letter lambda denoting wavelength, the distance between successive peaks of a wave.

ν —Greek letter nu, which denotes frequency.

σ —Greek letter sigma, used to denote the Stefan-Boltzmann constant in the Stefan-Boltzmann equation (Chapter 2).

Index

- abrupt climate change, 162–164
- absolute zero, 14
- acid rain, 104, 110, 175
- acidity, 126
- adaptive grid refinement, 158
- adiabatic system, 49
 - compression, 49–50
 - dry adiabat, 50
 - expansion, 49–50
- aerosols, 104, 142
 - sulfate aerosols, 143
- agriculture, 2, 92, 123–124, 153, 156, 162, 166, 181
- albedo, 20, 73–75
 - from clouds, 78–81
 - feedback, 73–75
 - of terrestrial planets, 23
- algae, 104
- alternatives for the future, 184–188
- altitude, pressure as a function of, 46–49
- anaerobic conditions, 105
- Antarctic ice sheet, 159
- Antarctica, 37, 75, 97, 119, 125, 131, 158–160, 163, 187
- anthropogenic climate forcings, 135
- anthropogenic interference, 181–184
- anticyclonic flow, 66
- Archimedes' principle, 159
- asymmetric stretch, 32
- Atlantic overturning circulation, 163–164
- atmosphere, 43–55
 - convection, 52–55
 - oxygen, 3, 29, 31–32, 51, 77, 89–92, 103–105, 119, 121, 129, 137, 149, 186
 - pressure as a function of altitude, 46–49
- atmospheric window, 32
 - optical thickness, 33
- averaging, 58–59
- band saturation effect, 34–38
- bare rock layer model, 19–25
- bending vibration, 30–31
- bias, 58, 135–137
- bicarbonate ion (HCO_3^-), 126
- biological climate impacts, 164–165
- biological pump, 128
- blackbody radiation, 9–16, 19–21
 - spectra, 15, 32
- boreholes, 146, 148
- breeder nuclear reactor, 185
- British Petroleum, 104, 106–107
- buffer chemistry, 127
- buildings, 180
- business-as-usual scenario, 178
- butterfly effect, 57
-
- CaCO_3 pH stat, 130
- cap and trade, 4, 174–175
- carbohydrates, 90
- carbon, 3–4, 84, 89–101, 119–131, 182
 - chemistry, 89–92, 97, 129
 - and climate, 3–4
 - fossil fuels, 3
 - inorganic, 90, 94, 126
 - juvenile, 99
 - offsets, 174–175
 - organic, 84, 90–97, 103, 120, 122, 125, 149
 - in sedimentary rocks, 89, 91–92, 97–99
 - sequestration, 180, 187
 - in soils, 92, 94
 - tax, 174–175
- carbon cycle, 3, 84, 119–137, *See also* perturbed carbon cycle
 - glacial cycles, 95, 97, 99
 - land biosphere, 3, 5, 89, 92, 103
 - the oceans, 94
 - weathering thermostat, 98
 - weathering, 98–99, 101, 121–122, 130–131
- carbon dioxide (CO_2), 2, 90, 122–131
 - CO_2 fertilization, 125
 - emissions, 5, 111–112, 115, 175, 182, 184
 - fertilization, 125, 128, 166
 - infrared activity, 43, 58

- carbon dioxide (CO₂) (*continued*)
 - missing sink, 124, 128–129
 - stabilization scenarios, 182, 184
- carbon on Earth, 89–101
 - reservoirs, 92
 - soil carbon pool, 92
 - terrestrial biosphere, 92
- carbonate ion (CO₃²⁻), 126
- carbon-cycle feedbacks, 84
- carbonic acid (H₂CO₃), 126
- chaos, 57–58, 144
- chemical bond, 30–31
- chemical equilibrium, 126
- chemical weathering, 98
- cirrus clouds, 78
- clathrates, 110
- clay minerals, 98
- climate, 57–70, *See also* forecasting climate
 - change
 - carbon, energy, and, 3–4
 - climate forcings, 4
 - fluctuating heat budget, 59–62
 - of last millennium, 145–147
 - borehole temperatures, 146
 - inverse methods, 146
 - proxy measurements, 145
 - variations on orbital time scales, 147–148
 - sensitivity parameter, 140
 - weather versus, 1–2
 - climate forcings, 135
 - anthropogenic, 135
 - changes in, 140–142
 - climate sensitivity parameter, 140
 - cosmogenic isotopes, 140
 - solar intensity, 140
 - volcanic eruptions, 142
- climate impacts, 153–172
 - abrupt climate change, 162–164
 - biological impacts, 164–165
 - biome types, 172
 - coral bleaching, 165
 - coral reefs, 165
 - dust bowl, 156
 - emergent effects, 155–158
 - human impacts, 165–166
 - ice and sea level, 158–162
 - polar bears, 165
 - precipitation changes, 170–171
 - rainfall, 156
 - temperature changes, 169
 - transpiration, 156
 - warming, 153–155
- clouds, 3, 78–81
 - cirrus clouds, 78
 - climate feedback, 3, 22, 69–70, 78–81
 - condensation nuclei, 80
 - Earth's energy balance, 26
 - cloud droplet size, 80
 - cumulus clouds, 78
 - stratus clouds, 78
- CO₂, *See* carbon dioxide (CO₂)
- coal, 103–104
 - sulfur in, 104
- compression, 49–50
- computer models, 57, 67, 69, 94, 129, 161
 - parameterization, 70
- conduction, 9
- Conference of Parties (COP), 177
- contrails, 80–81, 123
- convection, 43, 49, 52–55, 61, 75, 78, 155
 - of a compressible fluid, 53
 - dry, 54
 - moist, 54–55
- convectively unstable water column, 52
- COP, *See* Conference of Parties (COP)
- coral bleaching, 165
- coral reefs, 165
- Coriolis acceleration, 62–68
- cosmogenic isotopes, 140
- Cretaceous period, 100
- cumulus clouds, 78
- cyclonic flow, 66, *See also* hurricanes
- Dansgaard-Oeschger (D-O) events, 83, 163
- Day After Tomorrow scenario, 83, 164
- deep sea vents, 91
- deforestation, 123–124, 129
- degassing, 99
- dengue fever, 166
- dichlorodiphenyltrichloroethane (DDT), 90
- dipole moment, 31–32
- discount rate, 176
- dissolved inorganic carbon, 94
- Doppler shifting, 36
- dry adiabat in the atmosphere, 50, 54
- dry convection, 54
- Dust Bowl, 156
- Early Eocene optimum, 100
- Earth's energy budget, 79
- Earth-system climate sensitivity, 37
- eccentricity, 95–96
- ecological succession, 164
- economics, 174, 178–181
- Educational Global Climate Modeling (EdGCM), 69
- e-folding height, 48
- e-folding time scale, 48
- el Niño, 81–83, 146
- electric charge, 10

- electric field, 10–11, 30–32, 79
- electricity, 104–105, 184–186
- electromagnetic spectrum, 12
- electromagnetic waves, 10–13
- Electronic Numerical Integrator and Computer (ENIAC), 69
- electrons, 90
- emissivity, 14–15, 22
- energy, 9–16, 19–26, 29–39, 43, 46, 50–51, 54, 58–62, 68–70, 73–74, 77–91, 99, 103–117, *See also* fossil fuels
 - budget for the atmosphere, 24
 - budget for the Earth overall, 25
 - budget for the surface of the Earth, 24
 - and climate, 3–4
 - coal, 103–104
 - consumption, 111–117
 - energy intensity, 112
 - Kaya identity, 112
 - per capita, 111
 - efficiency, 180
 - hydroelectric, 103, 113, 115
 - intensity, 37, 112–115
 - nuclear, 103, 111, 184–185
 - oil, 106
 - solar, 4, 16, 18, 20–25, 30, 59–62, 70, 74, 76, 79, 95–97, 101, 103, 123, 140–143, 147, 155, 177, 185–186
 - sources, 103–111
 - supply, 180
 - wind, 5, 34, 57–58, 62, 66–70, 81–82, 103–105, 111, 124–128, 156–158, 185–186, 189
- ENSO, *See* el Niño
- enzymes, 91
- equilibrium vapor pressure, 51
- expansion, 49–50
- exponential function, 38, 47–48, 117
- external cost, 174

- Fast Ocean/Atmosphere Model (FOAM), 68
- FCCC, *See* Framework Convention on Climate Change (FCCC)
- feedbacks, 73–85
 - carbon-cycle feedbacks, 84
 - ice albedo feedback, 73–75, 128, 147, 154, 163
 - negative, 73
 - in paleoclimate record, 84
 - positive, 73
 - runaway, 75, 101
 - Stefan-Boltzmann feedback, 73
 - terrestrial biosphere feedbacks, 83–84
 - water-vapor feedbacks, 75–78
- Fennoscandian ice sheet, 95, 97
- fingerprinting, 177
- flood, 161, 166
- fluctuating heat budget, 59–62
- forecasting climate change, 2–3
- fossil fuels, 3, 89, 91–92, 97, 103–116, 124, 129, 183, 187, *See also* energy
- Foucault's pendulum, 63–64
- Framework Convention on Climate Change (FCCC), 4, 177
- freons, 119, 129
- frequencies, 10–13, 15, 30, 32–34, 36, 78–79, 95, 157–158

- gamma radiation, 11
- gases, 30–32, *See also* greenhouse
 - gases/greenhouse gas emission
- gasification of coal, 186
- Gaussian (bell) curves, 108
- GDP, *See* global domestic production (GDP)
- geoengineering, 187–188
- geopotential surface, 66
- geostrophic flow, 66, 70
- gigatons (Gton), 92, 104
- glacial cycles in atmospheric CO₂, 97
- glacial interglacial cycles, 84, 94, 131, 161, *See also* last glacial maximum
- glaciers, 1, 138–140, 157–160, 166
- global domestic production (GDP), 112–115, 180–181
- global temperature, 4
- global warming, 135
 - alternatives, 184–188
 - dangerous anthropogenic interference, 181–184
- economics, 178–181
- forecast, 3, 58, 70, 77, 144, 148
- economics of, 173–
 - cap-and-trade, 174
 - carbon offsets, 174
 - compound interest, 189
 - discount rate, 176
 - external cost, 174
 - market, 174
 - tax, carbon emission, 175
- negotiations, 176–178
- Great Red Spot of Jupiter, 66
- greenhouse effect, 4, 29–39, 155
 - from clouds, 80–81
 - effect of, 38–39
 - powered by lapse rate, 43–46
 - electrical dipole, 31–32
 - interaction with earthlight, 32–34
 - international negotiations, 4
 - layer model with, 23–26, *See also under* layer model
 - runaway, 75, 101

- Greenland, 37, 68, 83, 95, 97, 131, 153, 158–165
 - ice core, 75, 83–84, 97, 122–124, 140, 146–147, 162, 164
 - ice sheet, 2–3, 37, 83–84, 95, 97, 121, 127, 131, 146–147, 153, 156, 158–162
- Gulf Stream, 62, 66, 69, 83
- Hadley circulation, 77–78, 155–156
- Hansen, Jim, 181–182
- heat, 2–3, 9, 11, 19–22, 24–26, 43–46, 49–54, 58–62, 69–70, 75–77, 79, 83, 95–96, 99, 103, 105, 111, 120, 122, 135–137, 146–147, 155–158, 160, 163–166, 186
 - island effect, 135
 - transport, 62, 77, 110, 164
- Heinrich events, 160, 163
- Holocene, 95, 162
- Hoover Dam, 18
- hothouse climates, 161
- Hubbert's peak, 108–109
- human climate impacts, 165–166
- humankind and climate, 1–5
 - human-induced climate changes, 1
 - Little Ice Age, 2, 4
 - Medieval Climate Anomaly, 2, 4
 - risk assessment, 4–5
- humic acids, 89
- hurricanes, 157–158
- hydrocarbons, 90
- hydrogen, 3, 32, 77, 90, 101, 104–105, 110, 120–122, 126–127, 186
- hydrolysis, 186
- ice albedo feedback, 73–75, 128, 147, 154
- ice and sea level, 158–162
 - Antarctic ice sheet, 159
 - glaciers, 158
 - Greenland ice sheet, 159
 - ice sheets, 158
 - ice shelves, 158
 - ice streams, 160
 - Larsen ice shelf, 160
 - sea ice, 158
 - West Antarctic Ice Sheet, 159
- ice core data, 97
- ice sheets, 2, 37, 84, 97, 121, 131, 146–147, 153, 158
- ice shelves, 75, 158–160
- IGCC, *See* Integrated gasification combined cycle (IGCC)
- igneous rocks, 130
- inertia, 65
- infrared light (IR), 2–3, 11, 13, 22, 24, 30, 32–33, 36, 60
- inorganic carbon, 90
- integrated gasification combined cycle (IGCC), 180
- Integrated Science Assessment Model (ISAM), 115, 117
- intensity of infrared radiation, 14, 20
- Intergovernmental Panel on Climate Change (IPCC), 81, 115, 123, 136–137, 139, 142–143, 153, 158, 161, 176, 179
- inverse methods, 146
- ISAM, *See* Integrated Science Assessment Model (ISAM)
- Joules, 13, 103
- juvenile carbon, 99
- 8.2k event, 83, 162
- Kaya identity, 112–115, 154, 184
 - carbon efficiency, 115
 - energy intensity, 114–115
 - GDP per capita, 114
 - population and, 114
- Kelvins (K), 14
- kerogen, 89, 97
- kiloWatt (kW), 103
- kinetic energy, 29
- Kirchoff's law, 13
- kitchen sink analogy, 26, 73, 99–100, 121, 130
 - applied to atmospheric methane, 122
 - applied to heat energy, 58, 62, 120
 - applied to the ocean pH, 130
 - applied to atmospheric CO₂, 34, 36, 38, 80–81, 92–94, 97, 99, 124, 129, 143, 146, 148
- Kyoto Protocol, 5, 177–178
- la Niña, 81–83
- Lake Mead, 18
- land biosphere, 3, 5, 89, 91, 103
 - in carbon cycle, 131
 - climate feedback, 73, 81
- land breathes, 92–93
- lapse rate, greenhouse effect powered by, 43–46
- Larsen ice shelf, 160
- last glacial maximum, 2, 95, 147–148, 161
- latent heat, 50–51, 135
- Laurentide ice sheet, 95, 160, 163
- lava, 98
- layer model, 19–26
 - bare rock layer model, 19–23
 - with greenhouse effect, 23–26
 - energy budget for the Earth overall, 25
 - energy budget for the ground, 24
 - energy budget for the pane of glass, 24
- Le Chatelier's principle, 127
- light, 9–13, 30–32

- electric charge, 10
- electric field, 10
- frequencies, 10
- magnetic field, 10
- scattering, 79
- speed of light, 10
- wave number, 10
- wavelength, 10
- limestones, 97
- linear functions, 49, 58
- lipid, 91
- Little Ice Age 2, 4, 138, 147, 153–154
- long tail to the atmospheric CO₂ lifetime, 130
- magma, 98
- magnetic field, 10–11, 140
- malaria, 166
- market, 174
- Mars, 4, 20, 22, 25, 76, 89, 91, 101
- Maunder minimum, 2, 22, 140, 147, 154
- Medieval Climate Anomaly (+0.5°C), 2, 4, 147, 153
- melting ice on land, 158
- meridional overturning circulation, 83
- metamorphic decarbonation, 99
- metamorphism, 92
- methane, 2–3, 34, 37, 84, 90, 110, 120–122
 - infrared activity, 31
 - methane hydrates, 110–111
 - oxidation, 84, 89–91
 - sources of, 120, 122
- microwave sounding unit satellite, 137–138
- missing sink for CO₂, 124, 128, 129
- mitigation, 176
- mixing ratio, 29, 121
- moist adiabat in the atmosphere, 50, 54
- moist convection, 54–55
- mole, 126
- molecular vibrations, 30
- Montreal Protocol, 119, 178
- moon, 186
- Mt. Pinatubo, 142, 145, 188
- natural gas, 3, 90, 105, 110–111
- negative feedback, 73, 75, 80–81, 174
- nuclear energy, 103, 111, 184–185
- nuclear winter, 43
- obliquity of Earth's orbit, 96
- ocean acidification, 128
- ocean carbon cycle, 93–97
 - dissolved inorganic carbon, 94
 - glacial-interglacial cycles, 94
 - holocene, 94
 - last glacial maximum, 95
 - ocean carbon reservoir, 93
- ocean carbon sink, 125
- ocean circulation patterns, 81–83
 - ENSO, 81–83
 - meridional overturning circulation, 83
- ocean fertilization, 187
- ocean pH, 130
- ocean ventilation, 126
- OH radical, 120–122
- oil, 3, 90, 104–111, 115, 155, 187
 - peak oil, 107
 - reserves to production (R/P) ratio, 106
 - shales, 106
 - source rocks, 104–105
 - tar sands, 106
- orbital variations and climate, 95, 147–148
- orbitals, electronic, 30
- organic carbon, 90
- oxidation state, 89–91
- oxygen, *See* atmosphere: oxygen
- ozone, 3, 33, 44, 75, 119–120, 122, 178
 - as greenhouse gas, 120
 - hole, 119, 162
- Paleocene Eocene Thermal Maximum, 148–149
- paleoclimate record, feedbacks in, 84
- parameterization, 70
- partial pressures, 29, 51, 101
 - pCO₂, 29, 37, 93, 101, 115, 130, 182
- parts per million (ppm), 29, 92
- peak oil, 107, 109
- peat, 91, 104
- per capita energy consumption, 111
- permafrost, 84, 110, 125, 155, 164
- perturbed carbon cycle, 119–137, *See also* carbon dioxide (CO₂)
 - human-related climate drivers, 123
 - methane, 120–122
 - missing sink, 124
 - ocean carbon sink, 125
 - ozone, 119–120
 - smoking gun, 119
 - terrestrial carbon sink, 124
- PETM, *See* Paleocene Eocene Thermal Maximum (PETM)
- pH reactions, 126–128, 130, *See also* ocean pH
- phase diagram, 75
- photosynthesis, 3, 89, 91, 94
- photovoltaics, 186
- phytoplankton, 104, 128
- polar bears, 165
- polychlorinated biphenyls (PCBs), 90
- population growth, 47, 114, 124
- positive feedback, 73
- precession of Earth's orbit, 95–96

- pressure
 - broadening, 36
 - as a function of altitude, 46–49
 - gradient, 66
- proteins, 91
- proxy temperature records, 145–148, 156
- Qori Kalis glacier, 139
- radiative convective equilibrium, 54
- radiative equilibrium, 54
- radiative forcing, 80, 123, 140, 147
- radical chemistry, 54
- radioactive decay, 47–48
- rainfall, 156
- redox, *See* oxidation state
- reduction chemistry, 90
- reflected light, 20
- relative humidity, 51, 54, 77, 78
- reserves to production (R/P) ratio for oil, 106, 109
- respiration, 3, 91, 125
- Richardson, Lewis, 69
- risk assessment, 4–5
- rocks breathe, *See* solid Earth carbon cycle
- Rogner, Hans-Holger, 104
- rotation vector, 65
- runaway greenhouse effect, 75–77, 101
- saline aquifers, 187
- satellites, temperatures measured from, 137–138
- saturation vapor pressure, *See* relative humidity
- sea ice, 1, 75, 147, 153, 158–159, 164, 165
- sea level, 46, 48, 66, 92, 153–172
- sea surface temperature (SST), 81–82, 136–137, 157
 - measurement, 128, 136–137
- seasonal cycle, 58–59, 62, 93–94, 96, 145, 153
- seasons, 60–61, 93, 95–96, 128, 135
- sedimentary rock carbon pool, 97
- sedimentary rocks, 91
- sensible heat, 50, 135
- sequestration of CO₂, 187
- silicate rocks, 98–99
 - weathering CO₂ thermostat, 98–99, 131
- skin altitude, 45, 46
- skin temperature, 25, 43, 46
- slug theory, 183
- smoking gun, 119, 135–149, *See also*
 - thermometer records
- snakes, 16
- soil carbon pool, 91–92, 129
- solar cells/photovoltaics, 186
- solar energy, 20–21, 59, 79, 186
 - photovoltaic, 186
 - thermal, 186
- solar forcing of climate, 142, 147
- solar intensity, 140
- solid Earth carbon cycle, 97–101
 - chemical weathering, 98
 - clay minerals, 98
 - cretaceous period, 100
 - early Eocene optimum, 100
 - kerogen, 97
 - lava, 98
 - limestones, 97
 - magma, 98
 - sedimentary rock carbon pool, 97
 - silicate rock, 98
 - Urey reaction, 98
- Southern oscillation, 81
- spectrum, 13
- speed of light, 10–11
- SST, *See* sea surface temperature (SST)
- stabilization scenarios, 182, 184
- static stability, 52–53
- statically stable water column, 52
- Stefan-Boltzmann constant (σ), 14–15
- Stefan-Boltzmann equation, 13–14, 21, 33
- Stefan-Boltzmann feedback, 73
- storm surge, 155
- storm track, 155
- stratosphere, 4, 44, 77, 119–120, 138, 142, 145, 177–178, 187–188
- stratus clouds, 78
- stretching vibration, 30
- subsidence zone, 77, 155, 161
- sulfate aerosols, 80–81, 142–143, 187
 - albedo, 80, 140, 147
 - indirect effect, 80, 123
- sulfur, 80, 91, 104, 187
 - in coal, 103–104, 126, 175
 - in seawater, 91, 104, 122
- sunlight, 2–4, 15, 19–23, 26, 38, 60–62, 73, 79, 89, 91, 95, 97, 103–104, 120–121, 135, 164–165, 186–187
 - as climate driver, 142, 147
 - as power source, *See* solar energy
- sunspot cycle, 22, 140
- supercritical fluid, 101
- supersaturated water vapor, 5
- surface ocean mixed layer, 126
- tar sands, 106
- tax, carbon emission, 175
- temperature record, 1–2, 136–138, 142, 145–146, 148
 - global, 4
 - land, 137, 161
 - proxy, 145–146

- satellites, 137–138
- sea surface, 81–82, 136–137, 157
- terawatts (TW), 103, 184, 186
- terrestrial biosphere, 83–84, 92–94, 104, 124
 - carbon fluxes, 92, 129
 - drought feedback, 82–84
 - radiation, 15
 - transpiration, 83
- terrestrial carbon sink, 124
- terrestrial planets
 - albedos of, 23
 - temperatures of, 23
- terrestrial radiation, 15
- thermal expansion of water, 158
- thermal inertia, 61, 182
- thermocline, 81, 126
- thermometer records, 9, 135–137
 - from satellites, 137–138
 - urban heat island effect, 135
- tilt of Earth's axis of rotation, 60–61
- time step, 67
- tracers, 129
- Tragedy of the Commons, 4, 173–174
- transpiration, 83, 156
- tree rings, 145–146
- tropical cyclones, 157
- tropopause, 44, 46, 77
- troposphere, 32, 43, 77–78, 120, 142, 157, 185, 187
- turbulence, 69, 81, 126, 178
- turbulent cascade, 69
- typhoons, 157
- ultraviolet light (UV), 11, 77, 120
- undersaturated water vapor, 51
- uranium, 103, 185
- urban heat island effect, 135–136
- Urey reaction, 98
- vacuum, 9–10
 - electromagnetic waves carrying energy through, 10–13
- vegetation, 2, 83–84, 153, 156
- Venus, 4, 20, 25, 35–36, 75–77, 89, 91, 101
- vibrational energy, 11
- vibrations, 30–32
 - bending vibration, 30–31
 - stretching vibration, 30
- viscosity, 69
- volcanic CO₂ fluxes, 98–99
 - forcing of climate, 142, 177
- volcanic eruptions, 142
- warming, 153–155
- water vapor, 2, 50–51
 - in convection, 43, 49, 54–55, 61, 75, 78
 - infrared activity, 2, 11, 19, 25, 29–32
 - supersaturated, 5
 - undersaturated, 51
- water-vapor feedbacks, 75–78
 - Hadley circulation, 77–78
 - phase diagram, 75
 - runaway greenhouse effect, 75
- Watts, 13–14, 20–21, 103
 - kiloWatt (kW), 103
 - teraWatts, 103
- wave number, 10, 13–14
- wavelength, 10–11, 15, 19, 22, 36, 80, 137–138
- weather, 57–70
 - modeling, 65–70
 - versus climate, 1–2, *See also* climate
- weathering, 99–101, 121–122, 131
- wedges, 184–185
- West Antarctic Ice Sheet, 159
- wind energy, 185–186
- x-rays, 11
- Δt_x , 37–38, *See* climate sensitivity

